Essays in Behavioral Economics

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Felix Anand Chopra

aus Köln

2022

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Lorenz Götte
Zweitreferent:	Prof. Dr. Armin Falk
Tag der mündlichen Prüfung:	28. April 2022

Acknowledgements

First of all, I would like to thank my advisor Armin Falk. Over the past five years, Armin has greatly shaped my understanding of research and, most notably, economics. I am grateful for our many productive conversations, each of them leading to a deep insight, clarity, and exciting ways to approach the challenges in my work. In particular, I appreciate that I could always count on Armin's honest feedback, and I admire his unerring ability to identify the greatest strengths and weaknesses of each idea right-away. As I am writing this, I realize how difficult it is to adequately put into words how thankful I am for Armin's guidance and support. I could not have wished for a better advisor in any way.

I also want to express my gratitude to my advisor Lorenz Götte for his relentless encouragement and support to pursue the idea of my job market paper, his practical guidance on working with observational data, and the many great conversations we had along the way. I also want to thank Florian Zimmermann, who provided highly constructive feedback that improved the experimental design of every study I have conducted so far. I am also grateful to Chris Roth, who has not only been a fantastic coauthor and mentor, but has also become a dear friend of mine.

I consider myself extraordinarily lucky to have Peter Andre, Teodora Boneva, Philipp Eisenhauer, Armin Falk, Thomas Graeber, Ingar Haaland, and Chris Roth as my coauthors. I have learned a tremendous amount from each one of them and together we successfully overcame countless challenges. Peter and I started graduate school together, and over the years, he has been a trusted companion, a bright and inspiring sparring partner and, above all, a truly amazing friend.

I am also very fortunate to have shared the journey of graduate school with an outstanding group of peers, including Peter Andre, Ximeng Fang, Luca Henkel, Chui Yee Ho, and Tobias Rachidi, who made this an unforgettable and enjoyable experience. I would also like to thank the postdoctoral fellows at the *Institute for Behavior and Inequality*, Josh Dean, Leander Heldring, and Matt Lowe, for the many inspiring discussions about economics. I am grateful to the *Bonn Graduate School of Economics* for financial and administrative support.

I am immensely thankful to my parents, my brother, and my grandparents for their unwavering support over the years. Finally, I am most grateful to my partner and love, Moana. Nothing seems impossible with her by my side.

Contents

List of Figures

List of Tables

Introduction

Our decisions are shaped by our information and knowledge about the world around us. Mass media are an important source of such information. For example, people increasingly turn to popular mass media programs and trusted media celebrities for financial advice. Does information and persuasive communication from such media programs affect people's behaviors and attitudes? And if so, can policymakers leverage the broad reach of mass media for behavioral change interventions? People also use mass media to stay informed about politics, which is an important input for the functioning of democracies as informed voters are necessary to hold politicians accountable. Why do we observe media bias and misinformation in the market for news? And how can policymakers fight misinformation and steer consumers towards high-quality sources? These examples illustrate the importance of improving our empirical understanding of what drives the demand for information in media markets, and how, in turn, exposure to information from the mass media shapes economic behaviors.

Another important source of information are our own experiences and observations of the people in our community. For example, by observing climate-friendly actions such as reduced meat consumption, we learn about the prevalence of social norms and the normative expectations in our community. But do people correctly perceive social norms? And if not, can policymakers use information campaigns to correct these misperceptions? These questions are of particular relevance in the context of climate change, where the private nature of many climate-friendly decisions may lead to pessimism about others' efforts to fight climate change. At the same time, understanding people's motivation to fight climate change requires better knowledge of how, more generally, the intertemporal context of prosocial decisions affects people's willingness to behave prosocially.

This thesis consists of four independent research studies that broadly revolve around these questions. Each study draws on insights from psychology and economics, and utilizes methods from experimental economics. Moreover, each study draws part of its motivation from an applied policy question: How can we improve people's financial decisions (Chapter 1)? How can we fight misinformation and fake news (Chapter 2)? How can we increase support for pro-climate policies (Chapter 3)? And how can we promote prosocial behaviors in intertemporal contexts (Chapter 4)? Below, I briefly summarize each chapter.

Chapter 1: "Media Persuasion and Consumption: Evidence from the Dave Ramsey Show" investigates to what extent entertaining mass media programs can influence the primary economic decision of how much to consume. I provide evidence from the Dave Ramsey Show, which is the second most popular radio talk show in the US and dedicated specifically towards people who struggle financially. For over 30 years, the Dave Ramsey Show has argued that Americans overspend and under-save and encouraged Americans to change their behaviors. I exploit a quasi-natural experiment created by the staggered expansion of the radio show over a period of 15 years to estimate the causal impact of its message on household expenditures, which I obtain from a large household scanner panel. I complement this approach with a tailored survey experiment to shed light on the underlying mechanism. A set of three main findings emerge from my analysis. First, exposure to the radio show decreases monthly household expenditures. The effects are larger among households with initially high expenditures relative to their income. Second, households decrease their expenditures by purchasing fewer goods (extensive margin) rather than trying to pay less for their current basket of goods (intensive margin). Third, the experimental results suggest that the radio show affects behavior by changing people's attitudes towards spending and borrowing money. These findings demonstrate the potential of entertainment programs for interventions aimed at changing financial decisions and financial attitudes. Specifically, my evidence suggests that repeated messages on mass media about the value of savings and the cost of debt can encourage people to decrease their consumption.

Chapter 2: "Do People Demand Fact-Checked News? Evidence From U.S. Democrats" studies how fact-checking affects the demand for news, as measured by the decision to subscribe to a weekly politics newsletter in an online survey. The main treatment variation is whether the articles in the newsletter are fact-checked. On average, Democrats have a muted demand for fact-checking of a newsletter featuring ideologically aligned news (from *MSNBC*), even though fact-checking increases the perceived accuracy of the newsletter. However, this average effect masks substantial heterogeneity: Fact-checking decreases demand for politically aligned news among Democrats with strong ideological views and increases demand among ideologically moderate Democrats. Furthermore, fact-checking increases the demand for a newsletter with politically non-aligned news (from *Fox News*) for all Democrats irrespective of the strength of their ideological leanings.

These findings provide a proof of concept that non-instrumental motives can shape the demand for ideologically aligned news, which has relevance for models of media markets. In particular, the study's findings are inconsistent with theories in which consumers only care about the accuracy of news and point to the relevance of theories incorporating non-instrumental motives, such as a preference for belief confirmation. These findings also support demand-side explanations of media bias and misinformation in the market for news.

Chapter 3: "Fighting Climate Change: The Role of Norms, Preferences, and Moral Values" turns to the question of what drives people to take actions against climate change and support pro-climate policies. In a first step, this study documents that individual perceptions of social norms, economic preferences such as patience and altruism, as well as universal moral values predict people's willingness to fight climate change, as measured through an incentivized donation decision. In a second step, this study more closely examines perceived social norms, as they are malleable in the short term and can create a potent momentum either in favor of or against climate action. The study documents widespread misperceptions of social norms in the United States: People vastly underestimate the prevalence of climate-friendly behaviors and norms among their fellow citizens. Providing respondents with correct information causally raises individual willingness to fight climate change as well as individual support for policies aimed at mitigating climate change. The effects of the information intervention are strongest for individuals who are skeptical about the existence and threat of global warming, who are commonly difficult to reach but crucial for building up a broad alliance against climate change.

The observed underestimation of climate norms in the US can form a potent obstacle to climate action by potentially trapping Americans in an equilibrium with low climate engagement: Individuals are discouraged by the (mis)perceived lack of support, and they abstain from taking actions themselves, which sustains the pessimistic beliefs held by others – a phenomenon that has been dubbed pluralistic ignorance (Allport, 1924; Miller and McFarland, 1987). This study shows that a simple, easily scalable, and cost-effective intervention can correct these misperceptions and encourage climate-friendly behavior, suggesting that social norms should play a pivotal role in the policy response to climate change.

Fighting climate change can be conceptualized as a prosocial decision in an intertemporal context in that it involves a trade-off between costs for the self in the years to come and future benefits to humanity. **Chapter 4: "Intertemporal Altruism"** is motivated by the observation that, more generally, prosocial decisions typically involve consequences for the self and others that are spread out over time. Donations, for example, tend to create immediate costs to the donor and delayed benefits for others. However, the existing theoretical literature on prosocial preferences largely abstracts from the time dimension of utility flows. This study aims to fill this gap in three steps. First, it develops a conceptual distinction between *consequence-dated* and *choice-dated* prosocial utility in intertemporal contexts. If utility is consequence-dated, it accrues with a delay that corresponds to when the actual utility consequences for others materialize. If utility is choice-dated, it is realized in temporal proximity to the act of giving. Second, this study conduct a high-stakes donation experiment that comprehensively characterizes discounting behavior in self-other tradeoffs, which allows us identify different prosocial motives from their distinct time profile. The patterns in the data can only be explained by a combination of choice- and consequence-dated prosocial utility. A key finding is that people behave more prosocially when both the costs and benefits of the prosocial act are delayed further into the future, which cannot be rationalized with standard models of discounted utility. Third, this study quantifies the importance of choice-dated and consequence-dated prosocial utility in our setting by by estimating a parsimonious structural model.

Taken together, the chapters of this thesis illustrate the relevance of developing a better understanding of how people inform themselves, how they acquire information, and how this information shapes people's economic behaviors and beliefs in a variety of settings.

References

Allport, Floyd Henry. 1924. *Social Psychology*. Boston: Houghton Mifflin. [3]

Miller, Dale T., and Cathy McFarland. 1987. "Pluralistic Ignorance: When Similarity is Interpreted as Dissimilarity." *Journal of Personality and Social Psychology* 53 (2): 298–305. [3]

Chapter 1

Media Persuasion and Consumption: Evidence from the Dave Ramsey Show

Abstract: Can entertaining mass media programs influence individual consumption and savings decisions? I study this question by examining the impact of the *Dave Ramsey Show*, an iconic US radio talk show which encourages people to spend less and save more. To that end, I combine household-level expenditure records from a large scanner panel with fine-grained information about the geographic coverage of the radio show over time. Exploiting the quasi-natural experiment created by the staggered expansion of the radio show from 2004 to 2019, I find that exposure to the radio show decreases monthly household expenditures. This effect is driven by households with initially high expenditures relative to their income. In a mechanism experiment, I document that listening to the radio show has a persistent effect on people's attitudes towards consumption and debt. This suggests that attitudinal changes are a key mechanism driving behavioral change. My findings highlight the potential of entertaining mass media programs for interventions aimed at changing people's financial decisions.

Acknowledgements: I am grateful to Peter Andre, Alexander Cappelen, Stefano DellaVigna, Joshua Dean, Ruben Durante, Armin Falk, Lorenz Goette, Thomas Graeber, Matthias Heinz, Leander Heldring, Matt Lowe, Ulrike Malmendier, Moana Roepke, Christopher Roth, Bertil Tungodden, Johannes Wohlfart and Florian Zimmermann for helpful comments and discussions. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project B03) is gratefully acknowledged. Financial support from the Joachim Herz Foundation is also gratefully acknowledged. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1 – 390838866. Ethics approval for the experimental part was obtained from the German Association for Experimental Economic Research (#T7wapLjB, 07/20/2021). The main experiment was pre-registered at the AEA RCT Registry (#AEARCTR-0008050). *Disclaimer*: Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1.1 Introduction

Low savings rates and rising levels of household debt are a major problem in the US and many other countries.[1] Identifying effective policy responses has proven challenging (Gomes, Haliassos, and Ramadorai, 2021), rendering the question of how to promote responsible financial behaviors important. From a policy perspective, entertaining mass media programs may be particularly promising as they can reach a broad audience with a persuasive message. The private market already offers a variety of mass media content providing advice on how to make financial decisions. A relevant example is the *Dave Ramsey Show*—one of the most successful radio talk shows in the US—which argues that Americans spend too much and save too little. But do people act on what they are told? Are they capable of implementing the advice in practice? Or do they stay only for the entertainment?

In this paper, I examine the impact of the *Dave Ramsey Show* to tackle the following question. How and to what extent can entertaining mass media programs change the economic decision of how much to consume? I provide evidence from a quasi-natural experiment created by the staggered expansion of the *Dave Ramsey Show* over a period of 15 years. Using fine-grained variation in the geographic coverage of the radio show's broadcast over time, I document the radio show's impact on individual consumption levels in a large household scanner panel. Moreover, I combine a variety of tools ranging from text analysis of web-scraped audio records to machine learning methods to supplement my empirical analysis. To shed light on the behavioral mechanism, I conduct a tailored experiment examining the effect of the radio show on people's attitudes towards spending and borrowing money.

The *Dave Ramsey Show* is the second most popular radio program in the US with more than 20 million weekly listeners on over 600 affiliated radio stations. Each weekday, Dave Ramsey talks about personal finance and provides credit counseling for three hours. The radio show provides an attractive setting to examine mass media persuasion in the consumption and savings domain for three reasons. First, it explicitly aims to persuade its audience to change their behavior. For over 25 years, the *Dave Ramsey Show* has consistently broadcasted its key message that Americans spend too much and save too little. Dave Ramsey argues that Americans live beyond their means trying to keep up with the Joneses but fail to realize that "the Joneses are broke."[2] In the radio show, debt is portrayed as a symptom of conspicuous consumption and the negative consequences of debt are regularly highlighted.

1. For example, high household leverage has been linked to macroeconomic instability (Mian, Sufi, and Verner, 2017) and can impede people's ability to accumulate sufficient savings for retirement (Lusardi, Mitchell, and Oggero, 2020). Low savings rates may not always reflect optimal decisions but can instead result from behavioral barriers such as self-control problems (Laibson, 1997; Karlan, Ratan, and Zinman, 2014).

2. Dave Ramsey, *The Total Money Makeover*

Second, the staggered expansion of the radio show provides a source of quasi-natural variation in exposure to its message across time and space. Beginning in 1996, the radio show expanded to other media markets by licensing its content to local radio stations, averaging about one new station every other week over the next 25 years.

Third, the radio show has never changed its format and has consistently provided the same advice, which Dave Ramsey confirmed in an interview.[3] The effects of exposure to the radio show are thus comparable over time. To support this, I additionally examine the show's topics by analyzing all episodes uploaded on YouTube from 2013–2021. By using web-scraping to obtain the speech-to-text transcripts of these episodes, I circumvent the challenge that radio programs are not systematically recorded. A topic model estimated on a text corpus equivalent to about 3,000 hours of content suggests that the distribution of topics is stable over time.

My empirical strategy exploits the fact that the radio show was introduced in different media markets at different times to assess its impact on consumption. Specifically, I employ a difference-in-differences approach to estimate the causal impact of the radio show on consumption using variation in household-level expenditures before and after the local introduction of the radio show. The main identification assumption underlying this approach is that the timing of the radio show's introduction is unrelated to other factors driving household consumption. Anecdotal evidence from personal interviews with senior executives of the radio show suggests that the timing of the expansion was not driven by strategic considerations. Indeed, I find that the expansion is uncorrelated with baseline observables. As a more demanding test, I examine whether machine learning methods, which excel at uncovering non-linear statistical relationships, can predict the timing of market entry from observables. In a cross-validation exercise, I find that a random forest regression (Breiman, 2001) fails to predict the timing of the expansion from data about local economic conditions and the socioeconomic composition of the local population. Taken together, this evidence alleviates concerns about strategic entry.

To implement my empirical strategy, I combine comprehensive data on, (i), individual consumption and, (ii), the geographic coverage of the radio show's broadcast over time. In particular, I draw on 2004–2019 household-level scanner data from the Nielsen Homescan panel, which includes detailed information on the monthly grocery purchases of a large, geographically dispersed sample of US households. To determine the availability of the radio show, I collect novel data on the timing, technical specifications and geographic locations of its affiliated radio stations. I account for the influence of the topography and physical obstacles on radio signal strength by using a radio signal propagation model (Olken, 2009). This allows me to observe the staggered expansion of the radio show at the zip code-month level and identify when Nielsen households had access to the *Dave Ramsey Show*.

3. Interview with AllAccess (July 6, 2010)

I present three main findings. First, my main result is that exposure to the radio show decreases monthly household expenditures. The intent-to-treat effect on households living in areas that receive access to the radio show is a 1.3% decrease in expenditures. An event-study approach examining household expenditures up to twelve months before and after market entry confirms these findings and documents the absence of differences in pre-trends in expenditures, which supports the key identification assumption. Moreover, the event study shows that the impact of the radio show is stable and does not dissipate over the next twelve months. As individual exposure to the radio show is unobserved, I conduct a bounding exercise to better interpret the magnitude of the intent-to-treat effect. This exercise suggests that exposure to the radio show decreases household expenditures by at least 5.4%. From a policy perspective, not only the average effect of the radio show matters but also whether it persuades the intended target population. Examining heterogeneity in effects, I find that the decrease in expenditures is driven by households with initially high pre-exposure expenditures relative to their income, i.e., those who might benefit more from curbing their spending. In contrast, household income alone does not moderate the magnitude of the effect.

Second, I examine *how* households manage to decrease their expenditures. The answer to this question is not obvious because the radio show provides only limited guidance on this topic above and beyond its main advice to rigorously track and budget all household expenditures. In principle, households could choose to purchase less or try to pay less for their current basket of goods. I provide evidence that households decrease their expenditures primarily by decreasing the total number of products purchased. In contrast, I find economically insignificant effects on measures of frugal shopping behavior, such as purchasing products with a large package size or on-sale products that come at a discount.

Third, I study *why* households decrease their expenditures. A large part of the radio show is explicitly aimed at changing people's attitudes towards consumption and debt. A change in fundamental attitudes would explain the stability of the radio show's impact on behavior. I therefore investigate whether the radio show changes people's attitudes. As the observational data is limited to expenditure records, I conduct a pre-registered experiment with a representative sample of 1,500 Americans to address this question. In the main experiment, respondents are randomly assigned to a treatment group that listens to the *Dave Ramsey Show* and a control group that listens to a neutral audio recording. After respondents finish a module designed to obfuscate the study's purpose, I use items from validated scales to measure attitudes towards consumption (Richins and Dawson, 1992) and debt (Davies and Lea, 1995). I find that listening to the *Dave Ramsey Show* for a mere five minutes causes treated respondents to adopt more negative attitudes towards consumption and debt. For example, treated respondents have 24% of a standard deviation more negative attitudes towards conspicuous consumption. A robustness treatment shows that the effects are not driven by the choice of the audio recording used in the control group.

Despite the minimalist nature of the intervention, the treatment effects persist for at least a week as confirmed by an obfuscated follow-up survey, thereby allowing me to rule out experimenter demand effects (Haaland, Roth, and Wohlfart, forthcoming). A back-of-the-envelope calculation suggests that the change in attitudes may be sufficiently large to explain the magnitude of the decrease in expenditures documented in the scanner data.

My findings also hold under a series of additional robustness checks. For example, I replicate the decrease in expenditures using a more demanding empirical specification that only exploits residual variation in radio signal strength that can be attributed to the influence of physical obstacles on radio signals (Olken, 2009; Armand, Atwell, and Gomes, 2020). I implement this approach by controlling for the hypothetical signal strength that would be achieved in the absence of topographic obstructions in my main specification. Moreover, to alleviate concerns about biases in two-way fixed effects models caused by heterogeneous treatment effects over time (Goodman-Bacon, 2019; Callaway and Sant'Anna, 2020; Chaisemartin and D'Haultfœuille, 2020), I replicate the event study approach using the imputation estimator proposed by Borusyak, Jaravel, and Spiess (2021).

This paper makes several contributions to the literature. First, I contribute to the household finance literature by demonstrating the potential of mass media programs for behavioral interventions aimed at changing individual financial decisions. Specifically, my evidence from the *Dave Ramsey Show* suggests that repeated messages from mass media channels about the value of savings and the cost of debt can encourage people to decrease their consumption.[4] This suggests that delivering carefully designed messages through mass media could be an attractive complement to other behavioral interventions, such as providing financial education in order to raise financial literacy (Lusardi and Mitchell, 2007; Hastings, Madrian, and Skimmyhorn, 2013; Fernandes, Lynch, and Netemeyer, 2014).[5] This resonates with the findings from a nascent literature studying the effectiveness of edutainment interventions, i.e., a combination of education and entertainment, in developing countries (La Ferrara, Chong, and Duryea, 2012; Coville, Di Maro, Dunsch, and Zottel, 2019; Banerjee, La Ferrara, and Orozco-Olvera, 2020; Bjorvatn, Cappelen, Sekei, Sørensen, and Tungodden, 2020). For example, Berg and Zia (2017) find that fi-

4. My evidence also relates to research on relative consumption motives (Frank, 1985; Abel, 1990; Falk and Knell, 2004; Charles, Hurst, and Roussanov, 2009; Heffetz, 2011; Bursztyn, Ferman, Fiorin, Kanz, and Rao, 2018). I show that exposure to public messages criticizing the desire to "keep up with the Joneses" can make people less willing to spend.

5. Alternative approaches to encourage savings include, among others, changes in the choice architecture (Madrian and Shea, 2001; Carroll, Choi, Laibson, Madrian, and Metrick, 2009), peer influence (Beshears, Choi, Laibson, Madrian, and Milkman, 2015), or classical tax incentives (Chetty, Friedman, Leth-Petersen, Nielsen, and Olsen, 2014). See Beshears, Choi, Laibson, and Madrian (2018) for a comprehensive review.

nancial messages embedded in a South African soap opera encouraged people to borrow from formal banks rather than informal sources of credit.[6]

Second, more generally, this paper presents the first causal evidence that mass media programs can affect individual consumption *levels*. I thus contribute to the growing literature studying the social and economic impact of mass media by providing evidence of mass media persuasion in the core economic domain of consumption and savings decisions. Consumption and savings decisions differ conceptually from other domains where media persuasion has previously been documented, such as political behavior (Gentzkow, 2006; Della Vigna and Kaplan, 2007; Enikolopov, Petrova, and Zhuravskaya, 2011; Durante and Knight, 2012; Adena, Enikolopov, Petrova, Santarosa, and Zhuravskaya, 2015; Durante, Pinotti, and Tesei, 2019; Wang, 2021), violence and conflict (Dahl and DellaVigna, 2009; Della Vigna, Enikolopov, Mironova, Petrova, and Zhuravskaya, 2014; Yanagizawa-Drott, 2014; Armand, Atwell, and Gomes, 2020), or gender norms (Jensen and Oster, 2009; La Ferrara, Chong, and Duryea, 2012; Okuyama, 2019). Moreover, my findings suggest that mass media programs can affect people's materialistic orientation, consistent with the sociological perspective on mass media as a cultural agent of change (Hjarvard, 2008, 2013). While scholars have explored the relationship between mass media and *what* people consume, it has proven challenging to identify a causal effect of mass media messages on *how much* people consume. For example, in a related paper, Bursztyn and Cantoni (2016) carefully evaluate the impact of pre-reunification exposure to Western television on consumption choices in former East Germany. Interestingly, they find that advertisements in Western television affected what consumers purchased, but they find no effect on total expenditures. I thus shed light on the long-suggested influence of mass media on consumption levels (Belk and Pollay, 1985; Richins, 1987).

Third, I provide causal evidence that *non*-advertisement mass media content can influence people's consumption choices. The persuasive influence of mass media on consumer behavior has traditionally been the subject of research in the marketing sciences (see Bagwell, 2007, for a review). However, empirical research on advertisement mainly focuses on the effect on the sales of individual brands and firms rather than total household expenditures, with recent (meta-)studies suggesting that television (Lodish, Abraham, Kalmenson, Livelsberger, Lubetkin, et al., 1995; DellaVigna and Gentzkow, 2010; Shapiro, Hitsch, and Tuchman, 2021) and digital advertising (Blake, Nosko, and Tadelis, 2015; Lewis and Nguyen, 2015) are largely ineffective. Indeed, a key question since Marshall (1919) is whether advertisement

6. An important difference to edutainment interventions is that Dave Ramsey *explicitly* encourages people to change their behavior. In contrast, edutainment interventions rely on *implicit* persuasion in the sense that messages aimed at behavioral change are subtly embedded in the respective movie or soap opera, which has been theorized to lower barriers to behavioral change (Banerjee, La Ferrara, and Orozco-Olvera, 2020).

is "combative", resembling a tug-of-war between advertisers without affecting total expenditures (Chen, Joshi, Raju, and Zhang, 2009). My findings suggest that persuasive communication can, in principle, change total expenditures by shaping people's attitudes towards consumption.

More broadly, this paper relates to the literature studying the impact of charismatic individuals (Antonakis, Cianciolo, and Sternberg, 2004; Jones and Olken, 2005; Bassi and Rasul, 2017; Bursztyn, Rao, Roth, and Yanagizawa-Drott, 2020; Müller and Schwarz, 2020; Wang, 2021) and recent work on narratives in economics (Akerlof and Snower, 2016; Bénabou, Falk, and Tirole, 2020; Eliaz and Spiegler, 2020; Shiller, 2020; Schwartzstein and Sunderam, 2021). Dave Ramsey employs narratives of frugality and restraint (Shiller, 2020) and argues against what he perceives as a "consumerist culture." My evidence thus suggests that charismatic media personalities can use stories and narratives to change people's attitudes and behaviors.

1.2 Background

1.2.1 The Dave Ramsey Show

The *Dave Ramsey Show*, featuring its host Dave Ramsey, is one of the most successful US radio shows of the past decades and was ranked second place after Sean Hannity on *Talkers Magazine's* list of top radio talk shows in 2021.[7] About 20 million Americans tune in every week and as of 2021, 49% of Americans had heard of the radio show (YouGov, 2021b). Broadcasted from its studio in Nashville, Tennessee, the talk show airs Monday through Friday from 2–5 pm Eastern Time, which is the time of the day when radio consumption peaks.

Message. The *Dave Ramsey Show* talks about money, debt, and personal finance, with a focus on helping people to "get out of debt". This distinguishes it from other radio talk shows that—with the exception of only two other major consumer finance shows—exclusively discuss politics, culture, and sports.[8] The radio show has a dis-

7. Appendix Figure 1.B.2 shows consistently more Google searches for the radio show than for *Hannity*.

8. In 2020, there were only two other consumer finance radio talk shows among *Talker's Magazine* list of top 100 radio talk shows. The *Ric Edelman Show* provides investment advice and guidance on estate planning. In 2020, the radio show aired on 62 radio stations for two hours each Sunday. The *Clark Howard Show*, which stopped broadcasting in 2020, talks about consumer finances and provides advice on how to "spend less and save more", in particular by avoiding "scams and rip-offs." This radio show mainly provides tips on how to save money by making use of special deals, coupons, or one-off promotions, thus appealing to people who enjoy being frugal. However, it is less geared towards persuading people to change their behavior. Its audience size of 3.5 million weekly listeners is small compared to the *Dave Ramsey Show*, and only 29% of Americans had heard of the radio show in 2021 (YouGov, 2021a). Consumer finance programs on national television, such as *Making Money with Charley Payne*, mostly feature news about the stock market, discuss individual stocks and provide investment advice.

tinct and consistent message about consumption and debt: Americans live beyond their means trying to keep up with the Joneses, but fail to realize that the Joneses are “broke and living in debt, too.” Given this diagnosis, the radio show aims to persuade Americans to reduce their consumption:

> *“Financial peace isn’t the acquisition of stuff. It’s learning to live on less than you make, so you can give money back and have money to invest.”*
>
> *– Dave Ramsey*

Debt is consistently portrayed as a symptom of immature behavior, a failure of self-control, and a desire to impress others through conspicuous consumption:

> *“It is human nature to want it and want it now; it is also a sign of immaturity. Being willing to delay pleasure for a greater result is a sign of maturity. However, our culture teaches us to live for the now. ‘I want it’ we scream, and we can get it if we are willing to go into debt. Debt is a means to obtain the ‘I want its’ before we can afford them.”*
>
> *– Dave Ramsey, The Total Money Makeover*

The radio show thus uses an economic narrative based on Protestant values of frugality and restraint (Shiller, 2020). Appendix Section 1.E.1 provides further qualitative evidence documenting this narrative and an analysis of Dave Ramsey’s rhetoric can be found in Dori-Hacohen (2019).

The radio show additionally makes use of both positive and negative role models to support its main narrative. First, the radio show celebrates people who paid off their debt by having them explain how they achieved this goal before exclaiming: “I’m debt-free!” This ritual, called the debt-free scream, reinforces the idea that having zero debt is socially desirable.[9] Second, the radio show uses negative examples to explain its financial advice on how to cope with debt. Specifically, the main part of the radio show consists of live conversations between Dave Ramsey and people who called the studio line. After describing their financial situation and how debt has negatively affected their relationships or mental health, callers ask Dave Ramsey for advice. These calls reinforce the radio show’s philosophy that debt is harmful.[10]

Financial advice. The radio show promotes rules of thumb that foster habit formation and focuses less on teaching intricate financial concepts:

> *“Winning at money is 80 percent behavior and 20 percent head knowledge. What to do isn’t the problem; doing it is. Most of us know what to do, but we just don’t do it.”*
>
> *– Dave Ramsey, The Total Money Makeover*

For instance, the radio show recommends the “snowball” method of paying off debt, which involves paying the balances off in order of the smallest to the largest balance.

9. An example of a debt-free scream can be found here.

10. It is not unusual for Dave to be angry at the callers, call their behavior “stupid”, and provoke them: “When are you going to quit freaking spending money that you don’t have?”

While not minimizing total interest paid, immediate successes boost people's motivation (Brown and Lahey, 2015; Kettle, Trudel, Blanchard, and Häubl, 2016). Indeed, past research has shown that simple rules can often be more effective in promoting better financial outcomes (Drexler, Fischer, and Schoar, 2014). Similarly, the radio show advises people to set explicit budgets and plan all of their expenses ahead of each month to preempt overspending. People should then use one paper envelope per budget category and fill them with the corresponding cash amount. In order to become debt-free, the radio show recommends its step-by-step method called the "7 Baby Steps", which starts by saving $1,000 for an emergency fund to pay for unforeseen expenses. People should then apply the debt snowball to their non-mortgage debt before proceeding with the next steps. The show frequently discusses how to implement these steps in practice.

1.2.2 Program consistency

The radio show has made no major changes to the structure of its daily program, retaining a caller-driven format based on live conversations between Dave Ramsey and callers seeking advice. A key advantage of this setting is that the radio show provides consistent advice over time, which makes the experience of listening to the radio show in different time periods comparable:

> *"My advice never changes. My plan works in a good economy and a bad economy because it's all about getting control of your money."*
> *– Dave Ramsey in an interview with AllAccess.com (July 6, 2010)*

In this section, I provide additional suggestive evidence that the topics of these conversations remained similar over time. As radio talk shows are not systematically recorded (Sweeting, 2015), I obtain content data from YouTube via web-scraping (Kerkhof, 2020). Specifically, I use a Python script to obtain the speech-to-text transcripts and metadata of the 5,587 YouTube videos uploaded by the *Dave Ramsey Show* between August 13, 2013, and May 31, 2021. In total, these videos generated 647 million views and their transcripts capture around 3,000 hours of radio content.[11] As the radio show gradually started to use YouTube more over time, 94% of the data is from 2017 or later.

To shed light on the evolution of the topic distribution of the radio show over time, I use *Latent Dirichlet Allocation* (LDA, see Blei, Ng, and Jordan, 2003), which is a commonly used technique for topic analysis that aims to extract a fixed number of latent topics from unstructured text data (Gentzkow, Kelly, and Taddy, 2019). To apply this method, I partition the video transcripts into text documents containing the equivalent of five contiguous minutes of speech. I then train a LDA model

11. I apply a series of common processing steps to the raw text data, such as removing stop words and stemming words, which I discuss in more detail in Appendix Section 1.E.2.

with ten latent topics on this text corpus. The trained model then assigns to each document a probability distribution over topics.[12] Figure 1.2.1 displays topic shares from 2013–2021, obtained by averaging the predicted topic probabilities across documents. Reassuringly, the most common topics identified by the model capture central themes of the radio show: The largest topic is "financial problems" (22%), which refers to segments where callers describe their personal finances, how much debt they owe, and how debt has negatively affected their well-being and personal relationships. This topic thus reinforces the radio show's message that debt is harmful. This is followed by a topic capturing the celebration of people who paid off their debt by decreasing their consumption and standards of living during the so-called debt-free scream (20%). The least common topics are "Insurance" (3%) and "Health care" (3%).

Figure 1.2.1 provides suggestive evidence that there are no major trends in the topic composition over time. While topic shares can fluctuate across years, none of the topics is on a clear upward or downward trajectory. Moreover, in the 2017–2021 period which accounts for 94% of all uploaded content, the topic distribution is remarkably stable. This is not surprising given the industry wisdom that radio listeners expect program consistency (Perebinossoff, Gross, and Gross, 2005), which the vice president of Ramsey Media confirmed:

> *"Consistency in messaging is paramount. You must give the audience what they want and expect on a consistent basis."*
>
> *– Brian Mayfield in an interview with Inside Radio (January 25, 2019)*

These findings should, however, be taken with a grain of salt as material uploaded on YouTube is likely to be carefully selected to appeal to YouTube users, and may only offer a partial glimpse of what is discussed on-air. In particular, only a quarter of the text data comes from full-length episodes, which the radio show began to upload in 2019, while the remainder of the data comes from videos edited down to "highlights" of an episode. Restricting the topic analysis to full-length episodes—where scope for selection is more limited—reveals that the topic distribution is stable across years (as shown in Appendix Figure 1.E.3). Moreover, despite the COVID-19 shock, Figure 1.E.3 shows that the distribution of topics within full-length episodes changes little from 2019 to 2021.

1.3 Data

To study the impact of the *Dave Ramsey Show* on consumption, it is necessary to combine two types of data. First, one needs fine-grained information about the radio

12. Appendix Section 1.E.2 provides more details about the implementation. Appendix Figure 1.E.1 displays the topic-specific word distribution of the trained LDA model. Appendix Section 1.E.2.3 provides additional descriptive evidence from word frequencies and word co-occurrence rates

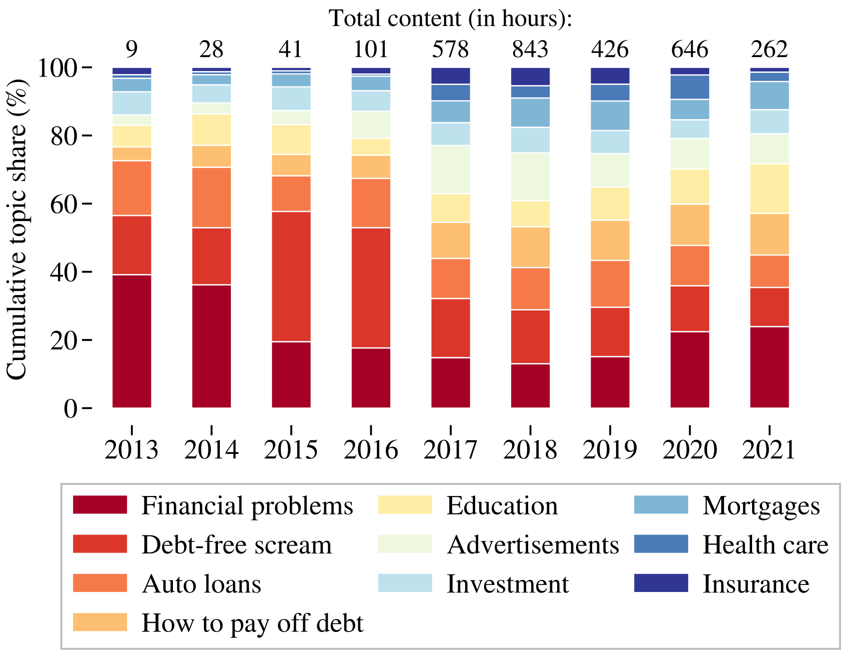

Figure 1.2.1. Topic distribution

Notes: This figure uses the text-to-speech transcripts of all videos uploaded on the *Dave Ramsey Show*'s YouTube channel between 2013 and 2021. The figure displays the distribution of topics across years. Topic shares are obtained from *Latent Dirichlet Allocation* by calculating the average probability of each topic across documents, where documents consist of 5-minutes of contiguous speech. For each year, the total content (in hours) uploaded on the radio show's YouTube channel is indicated above each bar. In total, excluding duplicate uploads, there are 2,934 hours of content.

coverage, and hence availability, of the *Dave Ramsey Show* across space and over time. Second, this information has to be linked with comprehensive, household-level expenditure records. This section describes the data and methods used to satisfy these requirements.

1.3.1 Radio coverage

As individual exposure to mass media programs is unobserved, I exploit variation in the *availability* of the *Dave Ramsey Show* across space and over time. To determine the availability of the *Dave Ramsey Show* at a fine-grained geographic level, I utilize a unique data set including information about the radio stations that broadcast the radio show at each point in time. I then determine the geographic coverage of these radio stations using an engineer-developed radio propagation model.

Ramsey Media provided a list of 493 radio stations that broadcast the *Dave Ramsey Show*, including their call sign, broadcasting frequency, and, crucially, the exact date they started carrying the radio show. As many radio stations build secondary transmitters to increase their service area, I manually match all listed radio stations with license and construction records from the Federal Communications Commission (FCC), which yields 176 additional secondary transmitters. Figure 1.B.3 displays the location of all 670 transmitters. For each transmitter, I collect technical specifications from the FCC's engineering records, such as the transmitter's effectively radiated power, height, broadcast frequency and geographic location, which I use to calculate the predicted receiver signal strength across zip codes.

The transmission of radio signals between a transmitter and a receiver location is governed by the laws of electromagnetic propagation. In free space, i.e., in the absence of topographic factors, radio signal strength depends on the frequency and power of the transmitter, and attenuates proportionally to the square of the distance from the transmitter. In practice, however, physical objects such as large buildings and topographic features such as mountains, forests and hills interfere with signal propagation, causing complex patterns of reflection, diffraction, and refraction (Cavell, Osenkowsky, Layer, Pizzi, and Hayes, 2017).

I therefore calculate the predicted radio signal strength corrected for topography using the Longley-Rice/Irregular Terrain Model (ITM).[13] Developed by the US government, the ITM is used by radio engineers and by economists, starting with Olken (2009), to predict the coverage area of radio transmitters. The high predictive accuracy of the model has been validated empirically in the field (Kasampalis, Lazaridis, Zaharis, Bizopoulos, Zettas, et al., 2013). Specifically, I calculate the path loss (in dB) between the transmitter location and the centroid of US zip codes. I then obtain the receiver signal strength by subtracting the path loss from the signal strength of the transmitter. Next, I use the maximum receiver signal strength across transmitters in a zip code to determine radio coverage (Durante, Pinotti, and Tesei, 2019). Finally, I combine the time-invariant geographic coverage of each transmitter with data on when these transmitters started to broadcast the radio show. The result is a monthly panel of the predicted receiver signal strength across zip codes between 1994 and 2019.

As radio coverage requires a sufficiently strong signal, I binarize the radio signal strength based on a threshold of 50 dBμV/m (Cavallo, 2017). This allows me to distinguish between zip codes with and without radio coverage in my analysis. In a validation exercise, I show that the results are robust to using thresholds between 40 and 50 dBμV/m (as shown in Appendix Figure 1.C.7), which have been used in prior work on the impact of radio broadcasts (Yanagizawa-Drott, 2014; Blouin and Mukand, 2019).

13. I thank Benjamin Olken for kindly sharing the ITM code.

1.3.2 Nielsen Homescan

To measure people's consumption, I draw on expenditure records from the Nielsen Homescan panel. A crucial advantage of this data compared to other household surveys is that the location of residency of each participating household is observed down to the 5-digit zip code level, which allows me to exploit fine-grained variation in radio coverage in my empirical analysis. The data set includes detailed information on the food and non-food product purchases of over 100,000 US households from 2004–2019. Households use an optical scanner at home to record information about their product purchases from grocery stores, drug stores, liquor stores and other retailers. The information includes the price, quantity, date of purchase, store identifiers, deals, and product characteristics at the Universal Product Code (UPC) level. For each shopping trip, households record the date and the store location before scanning the UPC bar codes of purchased items and entering prices and quantities. If the retailer exchanges point-of-sale data with Nielsen, the weighted average retailer-week price of each item is automatically recorded. Otherwise, households manually enter prices from their receipt and any deals involved in purchasing the item.[14] Nielsen imposes an undisclosed annual expenditure threshold that the value of all recorded purchases must exceed for a household to be included in the data set. Comparisons with the Consumer Expenditure Survey suggest that recorded purchases in the Nielsen panel account for a quarter of average annual household expenditures (Dubé, Hitsch, and Rossi, 2018). Nielsen also collects a broad set of self-reported demographic information, such as household income and household composition, age, gender, race, employment status and education of the household heads. Importantly, households also report the 5-digit zip code of their location of residency.

When recruiting panelists, Nielsen employs a stratified sampling approach to ensure that the sample is broadly representative of the general population in terms of nine demographic characteristics.[15] Moreover, the Nielsen Homescan sample is highly geographically dispersed. Figure 1.B.4 displays the distribution of Nielsen households across the 210 Designated Market Areas (DMAs), where a single DMA comprises several counties. These DMAs are used in the media industry to define media markets.

14. A potential concern is that households record product purchases with errors. Einav, Leibtag, and Nevo (2010) study the quality of the data by comparing scanner data from a large retailer with self-reported product purchases and find that the reporting error is comparable in magnitude to other commonly used economic data sets.

15. The demographic variables are household size, income, age of head of household, race, Hispanic origin, education of male and female household heads, occupation of head of household, presence of children, and county size. Lusk and Brooks (2011) study selection into household scanning panels such as Nielsen Homescan. They find that panelists tend to be older, more educated, more female and more price sensitive compared to a probability-based sample.

My primary outcome is the log of monthly household expenditures, which I obtain by aggregating total food and non-food expenditures before coupon use across all shopping trips within a calendar month. In Section 1.5.7, I verify that my results are robust to using alternative definitions of household expenditures. Appendix Figure 1.B.5 provides an overview of the geographic variation in average monthly household expenditures, which range from $357 up to $530 per month.

In my empirical analysis, I apply three exclusion criteria. First, I drop households that join the Nielsen Homescan panel after the *Dave Ramsey Show* became available in their location of residence. As my empirical strategy identifies the impact of the radio show from within-household changes, these "always treated" households do not contribute any identifying variation. On the contrary, recent progress on the econometrics of two-way fixed effects models shows that the presence of always treated units can actually bias estimates (see, for instance, Goodman-Bacon, 2019; Borusyak, Jaravel, and Spiess, 2021). Second, I focus on households that participate in the Nielsen panel for at least two years. For households that experience a change in radio coverage, I require that households are observed at least one year before and after they receive access to the radio show to ensure a sufficient observation period. Finally, I drop households that move across zip codes to address concerns about changes in purchase behavior in the years around the move (Bronnenberg, Dubé, and Gentzkow, 2012; Allcott, Diamond, Dubé, Handbury, Rahkovsky, et al., 2019). This additionally addresses concerns about selective migration of households into regions with access to the *Dave Ramsey Show*. The final panel of 3,744,078 household-months comprises 39,016 households in 11,219 zip codes across 202 DMAs.

1.3.3 Additional data

I supplement my analyses with additional data from various sources, including information on monthly house prices at the zip code level (from the Zillow Group), the county-level monthly unemployment rate and annual per-capita income (from the Bureau of Labor Statistics), the share of population in urban areas, racial composition and age groups (county-level; US Census and American Community Survey), and information about the Christian population. Moreover, I obtain county-level data on voter turnout and party vote shares for the 2000-2016 Presidential elections from the MIT Election Data and Science Lab (2018). Appendix Section 1.A provides an overview of all data sources.

1.4 Empirical strategy

1.4.1 National expansion

My empirical analysis exploits the staggered, national expansion of the *Dave Ramsey Show* across the US between 2004 and 2019. The radio show started in 1992 on 99.7 WWTN in Nashville, Tennessee, and began expanding to other markets in 1996. As a self-syndicated radio show, Ramsey Media neither owns nor operates radio stations, but rather engages in so-called affiliate relations with independent radio stations and networks. Affiliates receive locally exclusive access in exchange for advertisement minutes, a common practice in the radio industry that enables talk shows to realize economies of scale and radio stations to outsource the risk inherent in content production. As of 2019, the radio show is broadcasted by over 600 radio stations covering 208 out of 210 DMAs (see Appendix Figure 1.B.3). Figure 1.4.1 provides an overview of the staggered national expansion by indicating the biannual availability of the *Dave Ramsey Show* as well as changes in its coverage area. The expansion of the radio show into geographically distant media markets occurs early on and the sequence of the expansion does not appear to be driven by geographic considerations. Moreover, with about 40 new affiliates per year, the expansion was generally uniform over time (see Figure 1.B.1).

In line with these patterns, the radio show's expansion was not driven by strategic decisions. Qualitative evidence based on personal interviews with senior managers responsible for the expansion of the radio show's affiliate network suggests that the radio show did not prioritize media markets based on socioeconomic characteristics, trends in local economic outcomes, or consumer preferences. Instead, it focused on simply increasing the number of its affiliated stations:

> *"The main determining factor for choosing a market to enter is whether or not we are already on in that market. [...] We are either on or we're not on. And so even if we are adding a station and the listenership numbers are minimal, it's still better than zero and still better than not being on."*
>
> *– Personal interview with a senior manager*

The primary reason for this is that radio stations evaluate prospective talk shows based on their past performance in other markets. Indeed, these interviews reveal that radio stations often require evidence of successes in other markets before becoming an affiliate.[16] Moreover, it was important to document a growing number of affiliated radio stations in different regions, as some stations were concerned that the radio show might only find regional success. Thus, the radio show faced strong incentives to expand its network of affiliated stations in a variety of locations.[17]

16. Dave Ramsey in an interview with AllAccess.com (July 6, 2010).

17. It is a common practice in the radio industry to promote talk shows to hundreds of radio stations (Hendricks and Mims, 2018), which makes a targeted approach based on in-depth market

As a result, the timing of market entry was mainly driven by idiosyncratic demand for a non-political, general interest radio show that allows radio stations to diversify their program of predominantly political talk shows. After describing my empirical strategy, Section 1.4.3 will present additional, statistical tests suggesting that the timing of market entry was driven by idiosyncratic factors.

1.4.2 Econometric model

To estimate the effect of the *Dave Ramsey Show* on household expenditures, I employ a difference-in-differences strategy leveraging the radio show's staggered market entry across US zip codes from 2004–2019. Specifically, I estimate the following equation on a monthly panel of households:

$$\text{Outcome}_{itz} \;=\; \beta \text{Coverage}_{zt} + \phi_i + \psi_t + X'_{itz}\lambda + \varepsilon_{itz} \tag{1.4.1}$$

In the primary analysis, Outcome_{itz} is the log monthly expenditures of household i, residing in zip code z, at time t. Coverage_{zt} is a binary indicator variable taking value one if the *Dave Ramsey Show* is available in zip code z at time t and zero otherwise. In all specifications, I include household fixed effects, ϕ_i, and year-month fixed effects, ψ_t. The vector X_{itz} includes time-varying covariates that account for changes in the household's economic situation and local economic shocks. Household-level controls include log household income, household size, marriage status, employment status, and age indicators. Local economic conditions are proxied by zip code-level house prices and the local unemployment rate. In additional specifications, I further include state×year-month fixed effects or DMA×year-month fixed effects, which effectively restricts comparisons to within the same state or Nielsen media market. For inference, I use robust standard errors clustered at the zip code level at which the radio coverage indicator varies. The results are robust to using alternative clustering of standard errors.[18]

Equation (1.4.1) estimates the impact of the *Dave Ramsey Show* under the assumption that the timing of market entry is conditionally uncorrelated with pre-existing trends in household expenditures. Under this assumption, we can use changes in household expenditures in markets without radio coverage as a counterfactual for the evolution of expenditures in regions that receive access to the radio show.

To empirically evaluate the plausibility of this identification assumption, I present estimates from an event-study approach, which allows me to inspect the

research economically infeasible as establishing relationships takes time and is a labor-intensive process.

18. While Nielsen provides post-stratification weights, I do not weight households in my analyses because the set of households that experience a local market entry of the radio show is not nationally representative. For completeness, I show that the results are robust to using the Nielsen weights in Section 1.5.7.

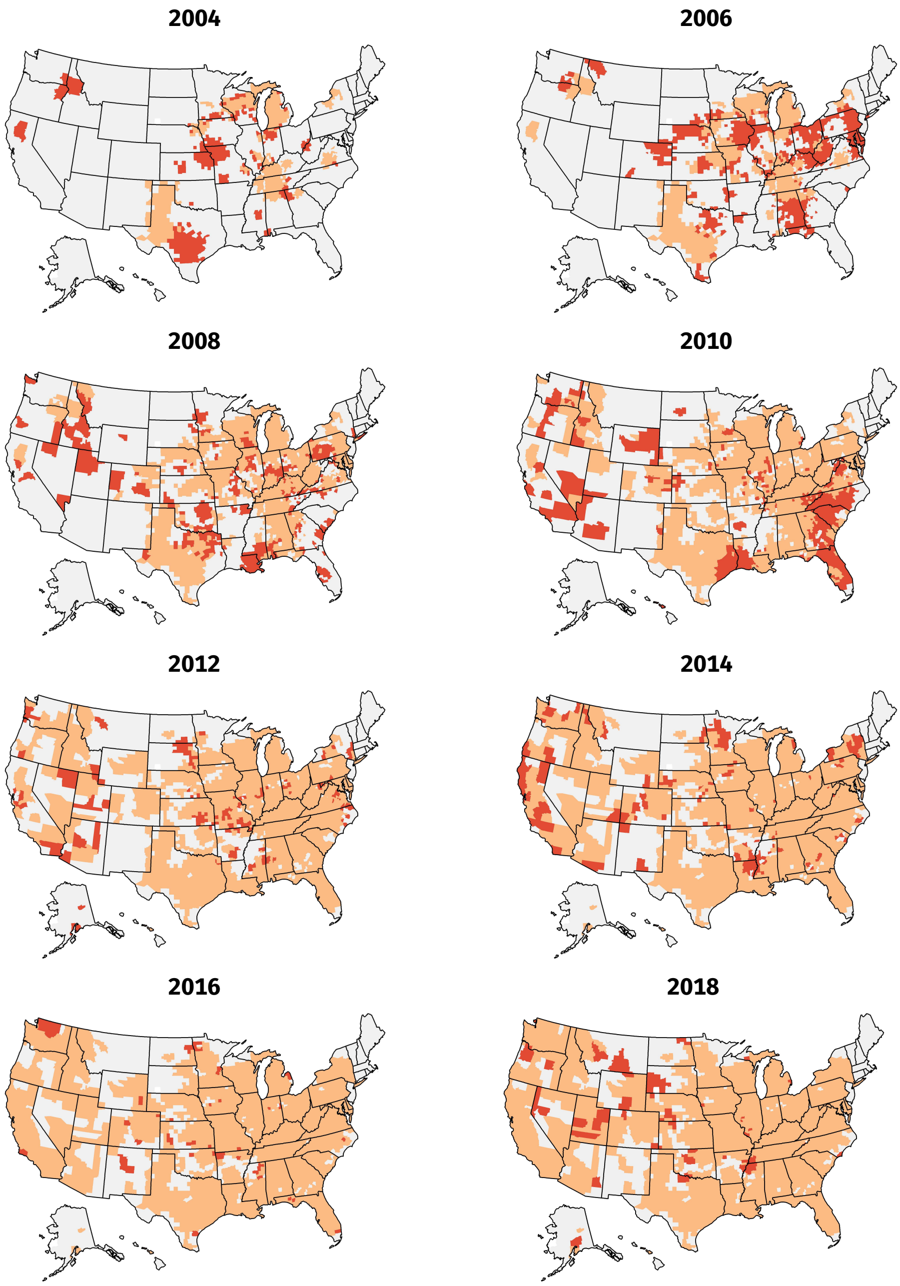

Figure 1.4.1. Radio coverage of the *Dave Ramsey Show*

Notes: This map shows the coverage of the *Dave Ramsey Show* from 2004–2018. Counties with coverage are shown in orange, while those without are indicated in grey. Areas that received coverage within the last two years are indicated in dark red. A county is defined as having coverage in this figure if at least 50% of the population has access to the radio show. The Longley-Rice/Irregular Terrain Model is used to estimate radio coverage at the zip code level, which I aggregate to the county level using population weights. Section 1.3 describes the data and procedure in more detail.

dynamics of short-term effects before and after market entry of the radio show. Specifically, I replace the binary coverage indicator in Equation (1.4.1) with a set of event-time indicators:

$$\text{Outcome}_{itz} = \sum_{\tau=-12}^{12} \beta_\tau \text{Coverage}_{zt\tau} + \phi_i + \psi_t + X'_{itz}\lambda + \varepsilon_{itz} \tag{1.4.2}$$

The binary event-time indicator $\text{Coverage}_{zt\tau}$ takes value one if $\tau = t - \tau_z^*$, where τ_z^* is the first time that the *Dave Ramsey Show* was available in zip code z, and zero otherwise. I further include binned indicator variables for event-times more than 12 months before and after market entry. After normalizing β_{-1} to zero, the coefficients β_τ capture the impact of the radio show τ months after market entry relative to the last month in the pre-exposure period. Given recent work on potential biases in two-way fixed effects models arising from dynamic treatment effects, I present complementary event-study estimates using the imputation estimator proposed by Borusyak, Jaravel, and Spiess (2021) as an additional robustness check in Section 1.5.7.

1.4.3 Identification assumption

This section provides further evidence supporting the plausibility of the identification assumption that the timing of market entry was conditionally exogenous.

1.4.3.1 Determinants of market entry

First, I examine the association between the first time the radio show became available in a given area and different standardized baseline covariates from the year 2000. As shown in Figure 1.4.2, these associations are all economically small and statistically insignificant. Specifically, a one standard deviation change in any baseline characteristic is associated with a change in the timing of market entry by no more than 3 months—a negligible association compared to the 39 months standard deviation of market entry. This complements the qualitative evidence that market entry was a non-strategic decision.

1.4.3.2 Machine learning

Second, I conduct a falsification test assessing whether one can predict the timing of market entry from sociodemographic factors. If observables do not improve the predictive accuracy, we should be less concerned about endogeneity of the staggered expansion. To provide a demanding test, I use supervised machine learning and cross-validation to assess the predictability of market entry. A key advantage of machine learning is that it can explore more general relationships and leverage higher-order interactions without imposing functional form assumptions such as linearity. In practice, I repeatedly train different models to predict the timing of market

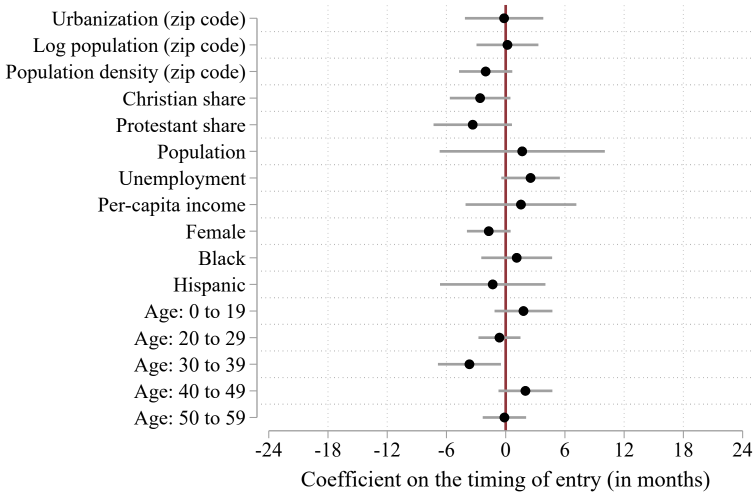

Figure 1.4.2. Determinants of the timing of market entry

Notes: This figure plots the coefficients from univariate regression of the first year-month of radio coverage on different baseline characteristics from the year 2000. The unit of observation are zip codes or counties depending on the level of aggregation at which baseline characteristics are measured. All baseline characteristics are standardized to have mean zero and standard deviation one to facilitate comparisons. The standard deviation of the timing of market entry is 39 months, or 3.25 years. 95% confidence intervals are constructed from robust standard errors clustered at the DMA level.

entry across zip codes with at least one Nielsen household between 2004 and 2018. I use the root mean squared prediction error (RMSE) on a hold-out sample to assess the model fit. The test-train sample splits are obtained from an implementation of a spatial leave-p-groups-out cross-validation approach to prevent "data leakage" from spatial autocorrelation (Le Rest, Pinaud, Monestiez, Chadoeuf, and Bretagnolle, 2014; Roberts, Bahn, Ciuti, Boyce, Elith, et al., 2017).[19] I then compare the distribution of the RMSE of each model to the distribution of the RMSE obtained from randomly assigning counterfactual entry dates.

Figure 1.4.3 presents the results. A "naïve" model making a constant prediction equal to the average entry date in the training data achieves a median RMSE of 3.9 years, with an associated p-value of 0.19 compared to the random benchmark distribution. Linear regression models using baseline observables do not improve the predictive accuracy ($p = 0.36$).[20] Next, I consider a Random Forest regressor (Breiman,

19. To split the data, I randomly draw three coordinates in the contiguous US and assign all zip codes within 500 km of these coordinates to the test data set. The training data comprises the complement after removing a "buffer zone" in the shape of a ring with a width of 300 km around the test data to ensure independence across samples. The diameter of the buffer zone was chosen such that the coverage area of a radio station does not intersect both the test and the training data.

20. The variables include the zip code and county population, population density, age shares (10-year bins), female, white, Hispanic and Christian population shares, per-capita income, the county unemployment rate and the degree of urbanization.

2001) with hyperparameters described in the Appendix, which is a commonly used general-purpose machine learning technique (Varian, 2014; Wager and Athey, 2018; Besley, Fetzer, and Mueller, 2019). Despite its flexibility, the Random Forest using baseline covariates has low predictive accuracy in this setting ($p = 0.49$). While this addresses endogeneity concerns based on baseline variables, the timing of market entry could also have depended dynamically on local trends in economic conditions. However, the results from a Random Forest regressor using panel data on the local unemployment rate and local average income suggest otherwise ($p = 0.47$). This evidence leaves little scope for local economic conditions to have driven the timing of market entry.

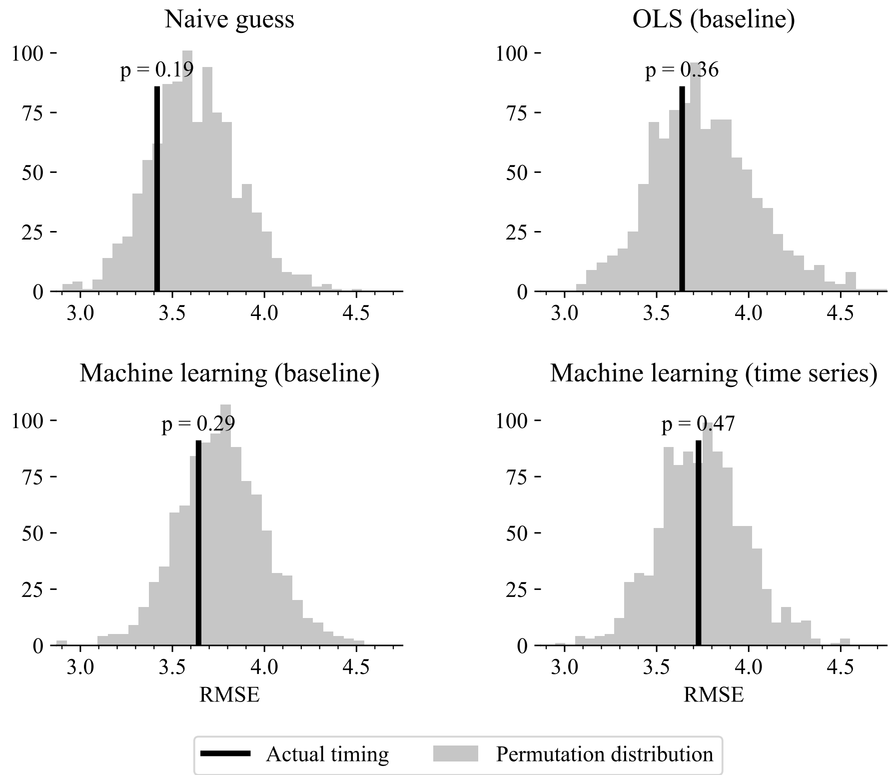

Figure 1.4.3. Predictability of the timing of market entry

Notes: This figure displays the results of a permutation test of the predictability of the timing of market entry across zip codes. In each panel, the black vertical line indicates the average root mean squared prediction error (RMSE) of the model obtained from a spatial cross-validation procedure. The implied p-values obtained from the permutation distribution are indicated. The permutation distribution is obtained from 1,000 random permutations of the dates at which affiliated radio stations started to carry the show, and subsequently recomputing the implied coverage across zip codes using the predicted signal strength. The "Naive guess" always predicts the empirical mean in the training data. "OLS (baseline)" and "Machine learning (baseline)" try to predict the timing of entry based on baseline zip code and county characteristics from 2000, including the demographic composition and local economic conditions. "Machine learning (time series)" shows the result of a Random Forest using annual data on the county unemployment rate and the average per-capita income from 2002–2016 as features.

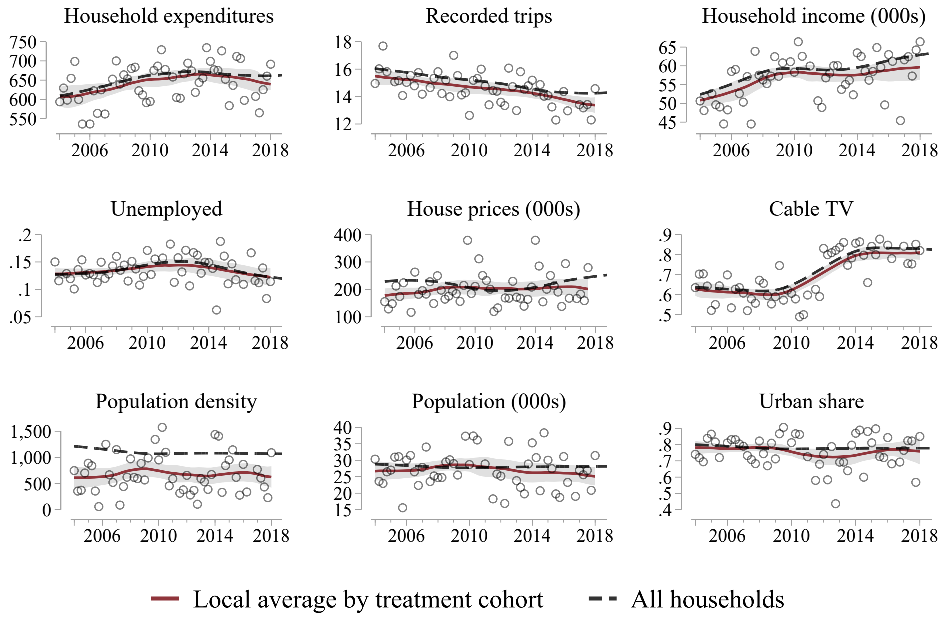

Figure 1.4.4. Covariates of households that gain access to the radio show

Notes: This figure compares the characteristics of households that received access to the *Dave Ramsey Show* in a given year-quarter to the average across all households in the Nielsen panel. The hollow circles indicate the quarter-by-quarter average characteristics of households that gained access to the *Dave Ramsey Show* for the first time in the given quarter. The red line indicates a smoothed local approximation of this average (Epanechnikov kernel, rule-of-thumb bandwidth estimator), with shaded areas indicating 95% confidence interval. Quarter-averages are trimmed at the 1st and 99th percentile prior to estimating the local approximation. The dashed black line indicates the local approximation of the quarter-by-quarter average characteristic of all Nielsen panelists. "Household expenditures" are total monthly expenditures in dollars. "Recorded trips" are the number of different shopping trips for which a household recorded purchases. "Household income" is measured annually. "Unemployed households" is the share of panelists that are unemployed at the beginning of the calendar year. "House prices" is the zip code-level Zillow House Price Index (monthly frequency). "Cable TV" is the share of panelists that have access to cable television. "Population density" is the zip code population density in 2010. "Population" is the zip code population in 2010. "Urban share" is the share of the population living in urban areas in 2000.

1.4.3.3 Characteristics of treatment cohorts

Idiosyncratic timing of market entry would imply that the covariates of incoming treatment cohorts evolve in parallel to the covariates of the average Nielsen household. If, however, the *Dave Ramsey Show* strategically timed its expansion based on information about the local audience, the characteristics of incoming treatment cohorts should change over time. To explore this, I group Nielsen households into different "treatment cohorts" based on the year-quarter in which they receive access to the radio show. For each treatment cohort, I then calculate the average covariates of these households in the year-quarter in which they are treated for the first time. Similarly, I calculate the average covariates of all Nielsen households for each year-quarter and subsequently compare differences in observables between incoming treatment cohorts and the average Nielsen household over time. Figure 1.4.4 presents the results. Each circle represents the average characteristic of the incoming treatment cohort, and the solid red line indicates a smoothed local average across treatment cohorts. The average across all Nielsen households is indicated by the black dashed line. The evidence suggests that incoming treatment cohorts are very similar to the average household at that point in time across a rich set of observables. For example, the expenditure levels of incoming treatment cohorts closely track average expenditures in the sample. This provides additional evidence suggesting that there was no selection based on observables such as household expenditures, income, local house price or population.

1.5 Results

1.5.1 Household expenditures

I first examine the impact of the *Dave Ramsey Show* on household expenditures. In Table 1.5.1, I estimate different versions of the baseline specification (equation 1.4.1) using the log of monthly household expenditures as the dependent variable. The main finding is that household expenditures decline after the market entry of the *Dave Ramsey Show*. Across specifications, I find a statistically significant decrease between 1.2% and 1.6%, which implies a decrease in annual expenditures of $70–93. Column 1 shows that when including only household and year-month fixed effects, the effect is a 1.3% decrease in household expenditures ($p < 0.01$). This effect remains statistically significant and quantitatively stable once I control for time-varying household characteristics ($p < 0.01$, column 2), which addresses concerns about household-level labor market shocks. Column 3 further controls for house prices and the local unemployment rate to account for heterogeneous trends in local economic conditions. The resulting decrease of 1.6% is slightly larger than the estimate without these controls. Moreover, the effect is robust both to the inclusion of state×year-month fixed effects that account for unobserved economic changes at the state level ($p < 0.01$, column 4), as well as to the inclusion of interactions between county baseline characteristics and year-month fixed effects (column 5).

Figure 1.5.1 presents the corresponding event-study estimates (equation 1.4.2) using log expenditures as dependent variable. The estimates show the absence of any statistically significant difference in pre-trends in the twelve months before market entry, supporting the plausibility of the identification assumption of parallel trends in household expenditures in the absence of the radio show. The effects are stable in the first year after market entry, suggesting a persistent change in behavior. The decrease in expenditures in the first months following market entry is consistent with the strong and immediate impact of listening to the radio show on consumption attitudes (see Section 1.6).

Table 1.5.1. Household expenditures

	Dependent variable: log (Expenditures)					
	(1)	(2)	(3)	(4)	(5)	(6)
Radio show	-0.0131***	-0.0128***	-0.0161***	-0.0121***	-0.0133***	-0.0140***
	(0.0027)	(0.0026)	(0.0027)	(0.0034)	(0.0034)	(0.0042)
N	3,744,066	3,744,066	3,407,700	3,407,700	3,355,677	3,354,689
R^2	0.518	0.521	0.522	0.524	0.525	0.529
Mean of dep. var.	6.185	6.185	6.186	6.186	6.185	6.185
Household & Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
Household controls		Yes	Yes	Yes	Yes	Yes
Local economic conditions			Yes	Yes	Yes	Yes
State x Time FEs				Yes	Yes	
County controls x Time FEs					Yes	Yes
DMA x Time FEs						Yes

Notes: This table uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable is the log of household expenditures. "Radio show" is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. Individual controls include the log of household income, age indicators, household size, married indicator and employment status indicators (full-time, part-time, unemployed). Local economic conditions comprise controls for house prices and the unemployment rate. Baseline county controls include the racial composition (share of whites), log per-capita income, log population and the share of Christians. Robust standard errors clustered by zip code are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

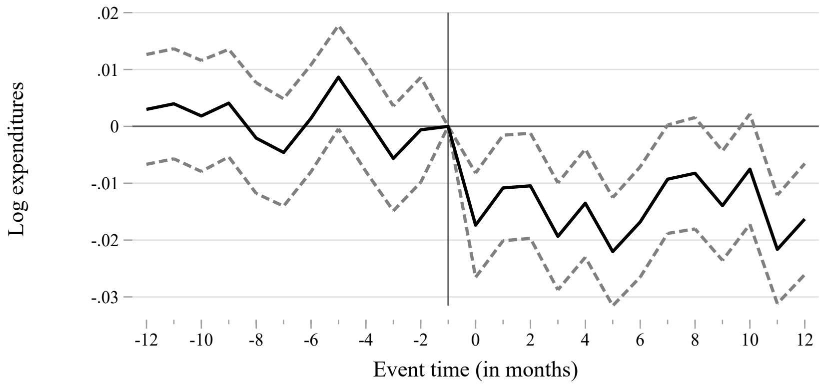

Figure 1.5.1. Event-study – Household expenditures

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level and depicts the results of a regression of log household expenditures on a set of event time indicators for the twelve months before and after market entry (see equation 1.4.2). The month before market entry serves as the omitted category. The regression also includes household and year-month fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level.

To address concerns that the decrease in expenditures reflects selection on unobserved economic shocks, I conduct two robustness checks. First, I estimate equation (1.4.1) on a subset of the data where the identification assumption is more likely to hold. Specifically, I exclude zip codes within 500 km of the radio show's headquarter in Nashville, Tennessee. Moreover, I exclude all 443 counties with an affiliated radio station. Appendix Table 1.C.6 shows that applying these exclusion criteria individually (columns 2 and 4) or jointly produces similar results (column 6). Second, I include DMA×year-month fixed effects in equation (1.4.1), which effectively restricts to local comparisons within the same media market. Trends in economic conditions are more likely to be comparable within media markets. Moreover, ratings information from market research companies are only available at the media market level. Flexible DMA-level trends thus account for selection on such unobserved information about audience preferences. In column 6 of Table 1.5.1, I find that the decrease in expenditures is robust to adding DMA×year-month fixed effects ($p < 0.01$). This evidence suggests that the results are not driven by endogenous market entry based on private information available to the radio show and its affiliated radio stations.

1.5.2 Magnitudes

As the empirical strategy identifies an intent-to-treat effect, the 1.3% decrease in expenditures is the most conservative estimate of the impact of the *Dave Ramsey Show* on the behavior of its audience. Ideally, one would use individual radio listenership information or geographically disaggregated audience data to estimate the local average treatment effect of the radio show's message on its actual audience. In the absence this data, I conduct a bounding exercise. Specifically, I divide the intent-to-treat effect by a range of alternative estimates of the share of Americans that have been exposed to the radio show's content, assuming that this percentage is constant across geographic areas.

An upper bound on the reach of the radio show are the 49% of Americans that have heard of the radio show (YouGov, 2021b), which implies a lower bound on the impact of the radio show on its audience of 2.7%. A lower bound on its audience can be derived from its weekly audience, which suggests a 6.5% national audience share. While this disregards sporadic and past exposure to the radio show, it implies an upper bound on the radio show's impact of about 20%. These bounds on the radio show's impact on its audience are likely to be non-binding, as they rely on very broad and very narrow notions of exposure. Alternatively, slightly tighter bounds can be derived from the following statistics. First, the radio show is "liked" by about 24% of Americans (YouGov, 2021b). Second, in my own representative survey, 8.3% of Americans can recall the name of the *Dave Ramsey Show* after listening to it for five minutes (see Section 1.6 for more details). These statistics would suggest that the *Dave Ramsey Show* causes a decrease in expenditures between 5.4% and 15.7% among its audience. The magnitude of the effect is thus economically meaningful, suggesting that mass media programs can have a substantial impact on the primary economic decision of how much to consume.

The magnitude of the effect is not implausible in light of the economically large impact of mass media on behavior documented in previous studies, in particular in settings where the media delivers an unusual message (DellaVigna and La Ferrara, 2015). These studies typically consider binary outcomes and calculate persuasion rates, a methodology pioneered by Della Vigna and Kaplan (2007), to compare media effects across settings. For example, Martin and Yurukoglu (2017) estimate that Fox News persuaded 58% of its viewers to vote Republican in 2000, while Wang (2021) finds that exposure to Father Coughlin's radio show persuaded 28% of his listeners to vote against Roosevelt. Moreover, Yanagizawa-Drott (2014) attributes 10% of the total violence during the Rwandan genocide to the impact of a popular radio station. While direct comparisons to voting or violent behavior are very difficult, the effect of the *Dave Ramsey Show* on consumption is consistent with the persuasiveness of mass media documented in other domains.

1.5.3 Purchased items

Next, I examine the mechanism through which households decreased their monthly expenditures. The radio itself provides comparatively little practical guidance on this question. Instead, its main advice is to "get on a budget" and keep track of all household expenditures to prevent overspending and impulse purchases. In light of this advice, one potential explanation for the decrease in expenditures is that households purchase fewer goods. To investigate this mechanism, I use the log of the total number of purchased items as a dependent variable, which I obtain by counting the number of UPC-level purchase records over the course of a calendar month. Table 1.5.2 provides estimates for different versions of the baseline specification (equation 1.4.1). I find that the availability of the *Dave Ramsey Show* causes households to purchase 1.7% fewer products ($p < 0.01$, column 1), which is robust across specifications (columns 2–6).[21] Figure 1.5.2 provides the corresponding event-study estimates, which indicate the absence of pre-existing differences in trends before the show's market entry. The implied effect of decreasing the number of purchased products on total household expenditures depends on the average price of the products which are no longer bought. Even if this price is 50% smaller than the price of the average product, a mechanism based on purchasing fewer goods would still account for at least half of the decrease in monthly household expenditures. This suggests that changes in the "extensive margin" are an important channel through which households decrease their expenditures.

1.5.4 Bulk and on-sale purchases

In addition to purchasing fewer products, it is ex-ante possible that households also try to reduce the amount they spend on their current basket of goods. Leveraging

21. Appendix Table 1.C.1 presents analogous estimates from a Poisson regression.

Table 1.5.2. Number of purchased items

	Dependent variable: log (Number of purchased products)					
	(1)	(2)	(3)	(4)	(5)	(6)
Radio show	-0.0168***	-0.0161***	-0.0210***	-0.0217***	-0.0232***	-0.0204***
	(0.0029)	(0.0028)	(0.0030)	(0.0036)	(0.0036)	(0.0046)
N	3,734,881	3,734,881	3,399,597	3,399,597	3,347,655	3,346,664
R^2	0.541	0.545	0.546	0.548	0.549	0.553
Mean of dep. var.	4.189	4.189	4.186	4.186	4.184	4.184
Household & Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
Household controls		Yes	Yes	Yes	Yes	Yes
Local economic conditions			Yes	Yes	Yes	Yes
State x Time FEs				Yes	Yes	
County controls x Time FEs					Yes	Yes
DMA x Time FEs						Yes

Notes: This table uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable is the log of the number of purchased items per month. "Radio show" is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. Individual controls include the log of household income, age indicators, household size, married indicator and employment status indicators (full-time, part-time, unemployed). Local economic conditions comprise controls for house prices and the unemployment rate. Baseline county controls include the racial composition (share of whites), log per-capita income, log population and the share of Christians. Robust standard errors clustered by zip code are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

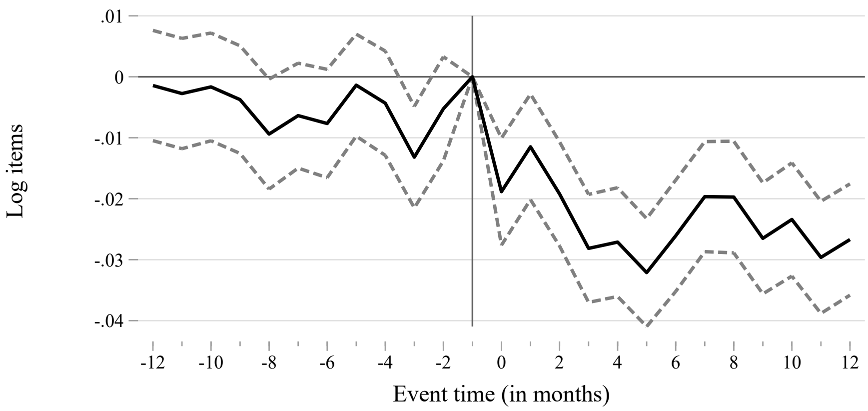

Figure 1.5.2. Event-study – Number of purchased items

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level and depicts the results of a regression of the log of the total number of purchased items per month on a set of event time indicators for the twelve months before and after market entry (see equation 1.4.2). The month before market entry serves as the omitted category. The regression also includes household and year-month fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level.

the richness of the Nielsen Homescan data, I construct two measures of savings efforts. First, I use UPC-level information about the packaging of each purchased product to construct a measure of bulk purchasing. Specifically, I rank products by their package size within their Nielsen product module. I subsequently construct the monthly share of expenditures accounted for by "large packages", which I define as belonging to the top quintile of the package size distribution. Second, using data on whether an item was purchased at a discount, I construct the expenditure share of discounted items. Table 1.5.3 reports the estimates of equation 1.4.1 using the measures of bulk purchases and discounted items as dependent variables. Columns 1 and 2 indicate more bulk purchases, as the share of expenditures accounted for by large items increases by approximately 0.5–0.6 percentage points ($p = 0.01$). Similarly, the expenditure share of on-sale items increased by about 0.3–0.4 percentage points ($p = 0.01$, columns 3–4). However, these ultimate effects of these behavioral changes on monthly household expenditures are likely to be modest compared to the effect of decreasing the number of products purchased, which is evident from the following example. Griffith, Leibtag, Leicester, and Nevo (2009) estimate a mean discount of 20% from purchasing on-sale items and average savings of 16% from purchasing bulkier products. Thus, the maximum decline in expenditures that is attributable to both activities is 15%, which suggests that the decrease in expenditures is primarily driven by the extensive margin.[22] Figure 1.5.3 presents the corresponding event study estimates for both measures.

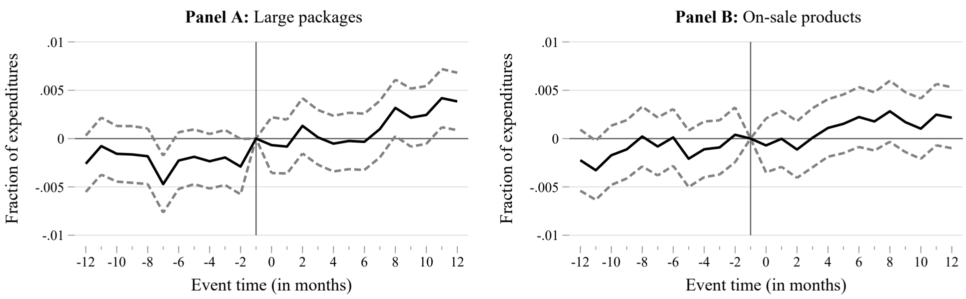

Figure 1.5.3. Event-study – Bulk purchases and on-sale products

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level and depicts the results of regressions of different outcomes on a set of event time indicators for the twelve months before and after market entry (see equation 1.4.2). The month before market entry serves as the omitted category. The regression also includes household and year-month fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panel A uses the share of expenditures accounted for by items in the top quintile of the package size distribution as dependent variable. Panel B uses the share of expenditures accounted for by items that were purchased on-sale as dependent variable.

22. The potential savings as a fraction of expenditures can be bounded from above by $0.006 \times 0.16 + 0.004 \times 0.20 = 0.00176$, which is 13.4% of the 1.31% decrease in overall expenditures.

Table 1.5.3. Bulk purchases and on-sale products

	Dependent variable: Expenditures share of					
	Large packages			On-sale products		
	(1)	(2)	(3)	(4)	(5)	(6)
Radio coverage	0.0043***	0.0047***	0.0064***	0.0035***	0.0029***	0.0043***
	(0.0007)	(0.0007)	(0.0008)	(0.0009)	(0.0010)	(0.0012)
N	3,734,872	3,399,588	3,399,588	3,734,881	3,399,597	3,399,597
R^2	0.460	0.463	0.465	0.714	0.714	0.716
Mean of dep. var.	0.290	0.290	0.290	0.299	0.305	0.305
Household & Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
Household controls		Yes	Yes		Yes	Yes
Local economic conditions		Yes	Yes		Yes	Yes
State x Time FEs			Yes			Yes

Notes: This table shows OLS regression estimates of equation (1.4.1) using a monthly panel of households. “Radio show” is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. The dependent variables are the share of monthly expenditures accounted for by purchasing items large items or on-sale products, respectively. “Large packages” is the share of expenditures accounted for by items in the top quintile of the package size distribution. “On-sale products” is the share of expenditures accounted for by items that were purchased on-sale. Robust standard errors clustered at the zip code are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.5.5 Heterogeneity

A heterogeneous impact of the radio show across different groups could be driven by (i) differential selection into the radio show’s audience, or (ii) differences in the susceptibility of these groups to the radio show’s persuasive messages. As individual exposure to the radio show is unobserved, it is difficult to distinguish these explanations, which makes it difficult to derive ex-ante hypotheses about which patterns of effects one would expect along dimensions such as gender, age, or education. However, from a policy perspective, it matters whether the radio show persuades the intended target population, i.e., those households that are likely to overspend. These households may both be more prone to listen to the radio show and more likely to follow its advice in light of the fact that the radio show is specifically geared towards people who “live beyond their means.” It is thus natural to hypothesize that initial expenditures moderate the magnitude of the radio show’s effect.

To test this hypothesis, I construct a proxy for baseline household expenditures and examine heterogeneity in effects along absolute and relative expenditures. Baseline expenditures are constructed as the average of inflation-adjusted monthly expenditures in the first year in the panel and excluding months in which households have access to the radio show.[23] I then separately estimate the impact of the radio show on expenditures among household whose baseline expenditures lie above or below the median household. Table 1.5.4 presents the results. I find a large and highly statistically significant effect among households with high baseline expenditures ($p < 0.01$, column 1). In contrast, column 2 reveals that the effect of the radio show is economically small and statistically insignificant among households with low baseline expenditures. Moreover, the negative point estimate in column 2 suggests that the effect among high-expenditure households is not driven by mean reversion. To examine whether this merely reflects differences in income, I construct baseline household income using the same procedure as above. Columns 3 and 4 show that there is no differential impact of the radio show among households with high or low baseline incomes. Indeed, when exploring heterogeneity by baseline expenditures relative to income, I again find a large decrease of 1.6% among households with high expenditures relative to their income ($p < 0.01$, column 5) and no statistically significant effect among households with low expenditures relative to their income (column 6). This evidence is consistent with the fact that the radio show's advice is geared towards people who overspend and suggests that the radio show primarily affects those who may stand to gain most from changing their behavior.

1.5.6 Exploiting topographic variation

This section considers a more demanding specification in which the impact of the radio show is identified using only residual variation in the continuous radio signal strength arising from the interaction between the timing of the staggered expansion and the influence of the local topography. This approach further alleviates endogeneity concerns based on strategic market entry as the factors driving market entry decisions are likely to be uncorrelated with local topographic variation. Specifically, I estimate the following equation:

$$\log\left(\text{Expenditures}\right)_{itz} \;=\; \beta \mathit{Signal}_{zt} + \gamma \mathit{SignalFree}_{zt} + \phi_{iz} + \psi_t + X'_{itz}\lambda + \varepsilon_{itz} \tag{1.5.1}$$

Signal is the standardized, continuous measure of signal strength in zip code z at time t, and *SignalFree* is its free-space analog, which differs from the former whenever topographic features interfere with the transmission of radio signals between the transmitter and the receiver location. By controlling for *SignalFree*, the main coefficient of interest, β, is only estimated from residual, plausibly exogenous variation in the radio signal strength. The nested zip code fixed effects account for any direct effects of topography on household expenditures. The identifying assumption underlying this approach is that the residual variation in signal strength arising from the

23. To account for the household composition, I normalize expenditures using an equivalence scale that assigns a weight of 0.7 to each additional adult and a weight of 0.5 to each child within a household.

Table 1.5.4. Heterogeneity analysis by expenditures

	Dependent variable: log (Expenditures)					
	Expenditures		Income		Expenditures / income	
	(1) High	(2) Low	(3) High	(4) Low	(5) High	(6) Low
Radio show	-0.019***	-0.002	-0.012**	-0.012**	-0.016***	-0.007
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
N	1,812,463	1,595,237	1,887,781	1,519,910	1,667,035	1,740,651
R^2	0.463	0.455	0.524	0.523	0.527	0.484
Mean of dep. var.	6.447	5.890	6.233	6.129	6.357	6.023
Full controls	Yes	Yes	Yes	Yes	Yes	Yes
State x Time FEs	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable is the log of household expenditures. "Radio show" is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. In all regressions, the set of control variables includes household covariates and controls for local economic conditions. Robust standard errors clustered by zip code are shown in parentheses. Each column provides estimates from a subset of households obtained by a median split based on the household covariate indicated in the column's header. For the median split in columns 1–2, I use the average, inflation-adjusted and equivalized expenditures in the first year a household is in the panel. For the median split in columns 3–4, I use the average, inflation-adjusted and equivalized household income in the first year in the panel. For the median split in columns 5–6, I use the average household expenditures normalized by income in the first year in the panel.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

interaction of the staggered expansion of the radio show and the detrimental effect of topographic obstructions on signal strength is uncorrelated with time-varying determinants of household expenditures. Appendix Figure 1.C.8 supports this assumption by documenting economically small and statistically insignificant correlations between the signal strength residuum and a large set of time-varying county-level characteristics.

Table 1.5.5 reports the results from estimating equation (1.5.1). Excluding the free space signal, a one standard deviation increase in signal strength leads to a statistically significant decrease in expenditures by 0.56% (column 1). Using only residual variation in radio signal strength, this effect increases to a 0.96% decline in expenditures per standard deviation change in signal strength (column 2). The effect is robust to including additional controls (columns 3–5). These estimates corroborate the baseline results and are quantitatively similar to the estimates from the specification using the binarized radio coverage variable presented in Table 1.5.1.

Table 1.5.5. Exploiting topographic variation in signal strength for identification

	Dependent variable: log (Expenditures)				
	(1)	(2)	(3)	(4)	(5)
Signal	-0.0056***	-0.0096***	-0.0088***	-0.0098***	-0.0082**
	(0.0016)	(0.0027)	(0.0027)	(0.0028)	(0.0037)
SignalFree		0.0049*	0.0039	0.0044	0.0092**
		(0.0028)	(0.0028)	(0.0029)	(0.0039)
N	3,599,959	3,599,959	3,599,959	3,272,490	3,272,490
R^2	0.521	0.521	0.524	0.525	0.527
Mean of dep. var.	6.185	6.185	6.185	6.186	6.186
Household & Time FEs	Yes	Yes	Yes	Yes	Yes
Household controls			Yes	Yes	Yes
Local economic conditions				Yes	Yes
State x Time FEs					Yes

Notes: This table presents OLS regression estimates of equation 1.5.1. "Signal" is the continuous measure of signal strength and "SignalFree" is the signal strength in free space. Both signal measures are standardized to have mean zero and standard deviation one. Robust standard errors clustered at the zip code level and shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.5.7 Additional analyses and robustness checks

Dynamic treatment effects. I conduct several robustness checks to address concerns arising from recent work on the econometrics of two-way fixed effects models (Goodman-Bacon, 2019; Callaway and Sant'Anna, 2020; Chaisemartin and D'Haultfœuille, 2020; Borusyak, Jaravel, and Spiess, 2021). Specifically, these studies show that two-way fixed effects estimators can be biased in the presence of heterogeneous treatment effects across cohorts and over time. First, I re-estimate equation 1.4.1, while excluding different treatment cohorts based on the year when they received access to the *Dave Ramsey Show*. Table 1.C.3 presents statistically significant estimates independent of which treatment cohorts are excluded. Notably, the results are robust to excluding households that receive access to the radio show during the Great Recession. Second, the Nielsen panel's sample is skewed towards top media markets as measured by population, which could bias results if across-market cohorts experience different dynamic effects. However, columns 1–4 of Appendix Tables 1.C.4 and 1.C.5 suggest that the effects are robust to excluding DMAs based on their Nielsen rank. Columns 5–7 show that the effects are additionally robust to focusing on homogeneous groups of markets, except for the lower tail where limited sample sizes become a concern. Third, I replicate the event-study using the imputation estimator proposed by Borusyak, Jaravel, and Spiess (2021), which is robust to dynamic treatment effects and efficient in finite samples. The estimates presented in Appendix Figure 1.C.1 closely resemble the dynamic patterns derived from OLS estimates (see Figure 1.5.1). This evidence suggests that treatment effect heterogeneity over time is not a major concern in this setting.

Falsification. To test whether the difference-in-differences estimates pick up spurious correlations, I conduct a falsification exercise. Specifically, I repeatedly assign a randomly chosen counterfactual market entry date to each zip code. If a zip code is outside the actual coverage area of all affiliated radio stations, the zip code is always assigned to the control group without any market entry. I thus vary only the timing of entry in a zip code, while the set of zip codes that receive access to the radio show remains unchanged. Based on the counterfactual timing of market entry, I apply equivalent sample restrictions (see Section 1.3.2) and re-estimate equation 1.4.1 using household expenditures as the dependent variable. Figure 1.D.1 compares the actual effect to the distribution of coefficients obtained from 500 repetitions of this procedure. The mean of the distribution is close to zero and negative, which reflects the fact that the set of treated zip codes is held constant and only the timing of entry changes. Moreover, the effect based on the actual timing of market entry is outside the empirical support of the distribution, suggesting that it is unlikely to arise by chance.

Placebo outcomes. As the *Dave Ramsey Show* is a non-political talk show, it should not affect political outcomes. I thus use the electoral turnout and the vote share of the Republican party as placebo outcomes. Specifically, I obtain county-level data for the 2000-2016 Presidential elections from the MIT Election Data and Science Lab (2018). Table 1.D.1 presents estimates from a panel regression of these political outcomes on the corresponding share of the county population that could listen to the radio show in the election year. As counties vastly differ in their population size, I weigh observations by the county's voting-age population. As expected, the radio show has no statistically significant effects on political outcomes.[24]

Additional robustness checks. The baseline results are robust to alternative specification choices. First, Appendix Tables 1.C.2, 1.C.9 and 1.C.10 document the robustness to (i) alternative constructions of household expenditures and the exclusion of outliers, (ii) alternative clustering of standard errors, (iii) using Nielsen's post-stratification weights. Second, Appendix Figures 1.C.2 and 1.C.5 document the robustness of the event-study approach to (i) the choice of control variables, (ii) state-specific trends, or (iii) replacing unit fixed effects with treatment cohort fixed effects (Imai and Kim, 2019). Third, the Nielsen Homescan sample is unbalanced for two reasons: Some households have missing purchase records for individual months, and households eventually leave the panel. While household fixed effects

24. It is difficult to construct a placebo variable using only data from the Nielsen Homescan panel because it is ex-ante not clear whether a particular product category should be unaffected by the impact of the radio show. For example, households could decrease their expenditures by using goods more efficiently and thus reducing waste.

already account for unobserved differences, compositional changes might affect the event-study estimates. I therefore re-estimate equation 1.4.2 on a balanced sample of households by excluding never treated households with gaps in their expenditure records and households that are not observed continuously during the event window. Despite reducing the sample size substantially, Appendix Figure 1.C.3 and 1.C.4 show that the results are robust to these changes. Fourth, Appendix Table 1.C.8 shows that the results are robust to excluding observations from the years following the introduction of the *Dave Ramsey Show* on other media channels, such as YouTube or satellite radio.

1.6 Experimental evidence

The above results reveal that the *Dave Ramsey Show* has economically large and meaningful effects on household behavior. The impact of the radio show is persistent and does not dissipate over the twelve months following market entry, which begs the question of how the radios how achieves persistent behavioral change. A distinguishing feature of the radio show is its regularly repeated narrative about consumption and debt—the notion that borrowing money and living beyond one's mean is wrong—which permeates every aspect of its three-hour program. Exposure to this narrative may cause people to revise fundamental attitudes towards consumption and debt, which would explain the persistence of behavioral change. While a multi-faceted radio program like the *Dave Ramsey Show* may also affect behavior through other channels, a mechanism based on attitudinal changes is likely to be particularly powerful. To examine the relevance of this mechanism, I conduct an experiment in which I exogenously vary whether respondents listen to the *Dave Ramsey Show* or a neutral audio recording before measuring attitudes. I provide evidence that listening to the *Dave Ramsey Show*'s narrative for *only* five minutes negatively affects people's attitudes towards consumption and debt.

1.6.1 Experimental design

1.6.1.1 Sample

The experiment was conducted in collaboration with *Lucid*, a professional survey company frequently used in social science research (Chopra, Haaland, and Roth, 2021; Haaland and Roth, 2021). To be eligible, respondents needed to reside in the US and be at least 18 years old. At the beginning of the survey, I screen out respondents that do not pass an attention check (see Appendix Figure 1.F.2). I also screen out respondents that cannot play audio files on their devices (see Appendix Figure 1.F.3), as this was a necessary technical requirement to administer the treatment manipulation. These exclusion criteria were preregistered (see Appendix 1.F.1). The final sample of 1,500 respondents is broadly representative of the general population in terms of age, gender, education and region (as shown in Table 1.F.1). Appendix Tables 1.F.2 and 1.F.3 present tests of balance to assess the integrity of the randomization procedure.

1.6.1.2 Main study

Panel A of Appendix Figure 1.F.1 provides an overview of experimental design. The full experimental instructions can be found in Appendix Section 1.F.4. The main experimental design was preregistered (see Appendix 1.F.1). Respondents first answer basic demographic questions and provide information about their personal finances. Then, respondents are randomly assigned to one of three experimental conditions: a treatment group, a control group, and a robustness control group.

Experimental conditions. The treatment group and the control group listen to different audio recordings, while the robustness control group proceeds without listening to anything.[25] The treatment group listens to a five minute audio recording of the *Dave Ramsey Show*, which was carefully chosen to include the major narrative elements of the show, such as the ubiquity of debt and the tendency of Americans to spend and borrow money to impress others. This allows me to mimic the experience of listening to the radio show for a longer period of time in which these elements would have naturally occurred. The control group listens to an unrelated podcast arguing that people should more carefully choose which "battles to fight" in their life. The podcast was deliberately chosen to hold many features constant, such as the total length, the gender of the speaker, the topical focus on self-help and personal improvement, and the narrator's paternalistic attitude. Appendix Section 1.F.5 contains a verbatim transcript of both audio recordings.

Obfuscation and delay. Experimenter demand effects induced by the audio recording might affect response behavior in the treatment group. To address this concern, I take several steps. First, I embed an obfuscation module directly after the audio recording. This module contains questions that mimic standard consumer research surveys, such as whether they would be more likely to listen to a particular radio station if it featured similar content. Second, I implement a "cool-off" period of about three minutes before measuring respondents' attitudes towards consumption and debt. Specifically, I elicit additional demographics, administer the "Big 5" financial literacy module (Hastings, Madrian, and Skimmyhorn, 2013), and measure demand for information about personal finances. Respondents should thus be uncertain about the primary interest of the study, which are attitudes towards consumption and debt—as specified in the pre-analysis plan.

25. Respondents cannot proceed to the next page for five minutes. They are told that they will have to answer some questions related to the audio recording after having finished listening to it, which serves to increase their engagement with the audio recording.

Outcome. To measure attitudes towards debt, I elicit respondents' agreement with four items from Davies and Lea's (1995) validated debt attitude scale. These items contain negative statements about debt, such as "There is no excuse for borrowing money." To measure attitudes towards consumption, I use two items from Richins and Dawson's (1992) validated materialism scale: "I admire people who own expensive homes, cars, and clothes" and "The things I own say a lot about how well I'm doing in life." Respondents' agreement with these items is measured on a 5-point Likert scale from "strongly agree" to "strongly disagree". For my primary analysis, I construct a (pro-)debt attitude index and a (pro-)consumption attitude index by summing responses to these items. Both indices are then z-scored using the control group mean and standard deviation. When estimating treatment effects on individual items, I recode answers such that larger values coincide with stronger agreement.

1.6.1.3 Follow-up survey

To shed light on the persistence of treatment effects over time, I conduct an obfuscated follow-up survey exactly one-week after the main experiment without administering any additional experimental treatments. I obfuscate the link between the main experiment and the follow-up survey by using a different survey layout and consent form, again eliciting basic demographics and including an additional obfuscation module measuring people's satisfaction with their primary bank. I then re-elicit attitudes towards consumption and debt using the original instructions from the main experiment. I managed to recontact 522 respondents, which corresponds to a recontact rate of 35%. Appendix Table 1.F.9 documents balanced baseline covariates in the follow-up survey, and Appendix Table 1.F.8 shows that there is no differential attrition across treatment arms.

1.6.2 Experimental results

Treatment effects. Table 1.6.1 documents the main result that the treatment effect of listening to the *Dave Ramsey Show* for five minutes causally affects people's attitudes towards consumption and debt.[26] In the main experiment, treated respondents have 53% of a standard deviation more negative attitudes towards debt and borrowing money compared to respondents in the control group ($p < 0.01$, column 1). They also have 24% of a standard deviation more negative attitudes towards consumption ($p < 0.01$, column 2). The magnitudes of the treatment effects are economically large and suggest that narratives embedded in mass media programs

26. I include the numerical age and age squared, log income, female indicator, and an indicator for having completed a Bachelor's degree or higher as control variables. Table 1.F.5 provides estimates without control variables. Appendix Table 1.F.6 shows that treatment effects on attitudes are not driven by individual items used to construct the indices.

have the power to substantially affect people's attitudes. The effects are robust to using respondents who did not listen to an audio recording as a comparison group (columns 3–4), suggesting that the treatment effects are not an artifact of the audio recording used in the control group.

In the obfuscated one-week follow-up survey, I find that treated respondents still hold 30% of a standard deviation more negative attitudes towards debt compared to control group respondents ($p < 0.01$, column 5).[27] This corresponds to 57% of the original effect size. Similarly, I still find a negative effect of 21% of a standard deviation on consumption attitudes ($p < 0.05$, column 6), which is an economically large effect in light of the minimalist intervention of listening to the radio show for a mere five minutes in the previous week.[28] Appendix Section 1.F.3 provides additional results from secondary outcomes suggesting that the effect of the radio show is driven primarily by changes in attitudes.

Table 1.6.1. Treatment effects on attitudes across studies

	Main study		Robustness: Passive control		One-week follow-up	
	(1) Debt attitudes	(2) Consumption attitudes	(3) Debt attitudes	(4) Consumption attitudes	(5) Debt attitudes	(6) Consumption attitudes
Treatment	-0.530***	-0.237***	-0.603***	-0.230***	-0.303***	-0.208**
	(0.065)	(0.061)	(0.061)	(0.060)	(0.094)	(0.090)
N	962	962	1,030	1,030	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the dependent variables are attitudes towards consumption and debt. The debt attitude index and the consumption attitude index are constructed as described in the main text and oriented such that larger values correspond to more positive attitudes towards the object. Both indices are normalized to have mean zero and standard deviation one. "Treatment" is a binary indicator taking value one for respondents who listened to a five minute recording from the *Dave Ramsey Show*. Columns 1 and 2 use data from the main experiment, focusing on the subset of respondents assigned to the treatment group and the control group. Columns 3 and 4 focus on respondents from the main experiment that were assigned to the treatment group or the robustness control group. Columns 5 and 6 use data from the one-week follow-up survey and pools respondents from both control group conditions (neutral podcast and no audio) as a joint control group. Control variables include numerical age and age squared, log income, female indicator, an indicator for having completed a Bachelor's degree or higher, and region indicators. Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

27. In the follow-up survey, I pool recontacted respondents from both control groups to maximize statistical power. I obtain quantitatively similar effect sizes without pooling these experimental groups. The results are robust to using inverse probability of attrition weights obtained from regressing a binary attrition indicator on a comprehensive set of baseline covariates (as shown in Appendix Table 1.F.7).

28. The test-retest correlation of the attitudinal measures is high with 0.60 and 0.74 for the debt and consumption indices. respectively.

Attitudes and behavior. The experimental findings raise a question about downstream effects of attitudinal changes on people's behavior, and in particular whether the treatment effect on consumption attitudes is large enough to explain the decrease in expenditures observed in the Nielsen panel. To address this question, I utilize correlational evidence from respondents in the control groups. Table 1.F.4 shows that consumption and debt attitudes correlate with self-reported behavior. Specifically, having a standard deviation more positive attitudes towards debt is associated with a 39% increase in personal debt ($p < 0.05$, column 1) and a 4.5 percentage points lower probability of having no debt ($p < 0.01$, column 3). Consumption attitudes are associated with a 17% increase in past spending on food, restaurants and leisure activities per standard deviation ($p < 0.01$, column 5). These correlations are robust to including sociodemographic controls (columns 2, 4, 6).[29] Assuming that the correlation of 0.1 between consumption attitudes and past expenditures in column 6 of Appendix Table 1.F.4 is causal, the 20.8% of a standard deviation decrease in consumption attitudes in the follow-up survey would imply a decrease in expenditures of about 2%, which is in the ballpark range of the observed effect of 1.3–1.6%. This back-of-the-envelope calculation provides additional support for changes in attitudes as a key mechanism through which the *Dave Ramsey Show* affects household behavior.

1.7 Concluding remarks

This paper provides causal evidence of mass media persuasion in the core economic domain of consumption and savings decisions. Specifically, I show that exposure to a popular US radio talk show arguing that Americans overspend and under-save causes people to decrease their consumption. To identify the causal impact of the radio show, I exploit quasi-natural variation in the availability of the radio show created by its staggered expansion from 2004 to 2019.

I provide three main results. First, I document that exposure to the radio show decreases household expenditures. Event-study estimates suggest that the effect of the radio show is not short-lived and instead persists for at least one year after the local introduction of the show. Second, I examine how households decrease their expenditures. My evidence suggests that the decrease in expenditures is best explained by households purchasing fewer products rather than exerting more effort to decrease the price of their current basket of goods. Third, I shed light on the underlying mechanism using a pre-registered experiment. I find that exposure to the radio show's message for a mere five minutes has an economically large and persistent, negative effect on people's attitudes towards consumption and debt.

29. Reassuringly, the attitudinal measures capture conceptually distinct facets: Consumption attitudes do not correlate with debt conditional debt attitudes (columns 1 and 3 of Appendix Table 1.F.4), while debt attitudes do not correlate with spending conditional on consumption attitudes (column 5).

My findings inform the debate on which policies are likely to be effective in mobilizing savings efforts. The evidence from the *Dave Ramsey Show* suggests that people act on the financial advice provided by mass media programs. Specifically, households are responsive to repeated messages on mass media advocating savings behaviors and cautioning against household debt. The finding that the radio show has larger effects among households with initially high expenditures relative to their income further suggests that the *Dave Ramsey Show* might have had positive effects from a welfare perspective.

This suggests that entertaining mass media are a promising avenue for behavioral change interventions aimed at improving financial outcomes. Financial advice on entertaining mass media programs, such as the *Dave Ramsey Show*, can reach millions of people on a regular basis at comparatively low marginal cost compared to other approaches such as classroom-based financial education programs. Moreover, entertaining mass media programs may appeal to people that are otherwise difficult to reach because of lacking interest in household finance.

However, effectively leveraging the power of mass media for behavioral interventions is not without its own limitations. For instance, it requires access to and collaboration with media production firms to tap their knowledge on how to design a product that is entertaining enough to appeal to a broad audience, while at the same time including carefully crafted messages aimed at behavioral change. This naturally constrains the type of information that can be disseminate through mass media. Whereas other channels might be better suited to teach intricate and detailed financial concepts, my evidence suggests that mass media can be used to raise awareness and change people's attitudes towards important issues such as insufficient retirement savings. Mass media interventions are hence best utilized in concert with a broader mix of policies and interventions aimed at improving financial outcomes.

References

Abel, Andrew B. 1990. "Asset Prices Under Habit Formation And Catching Up With The Joneses." *American Economic Review* 80 (2): 38–42. URL: https://www.jstor.org/stable/2006539%20http://www.sciencemag.org/cgi/doi/10.1126/science.151.3712.867-a. [9]

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of The Nazis in Prewar Germany." *Quarterly Journal of Economics* 130 (4): 1885–939. DOI: 10.1093/qje/qjv030.Advance. [10]

Akerlof, George A., and Dennis J. Snower. 2016. "Bread and bullets." *Journal of Economic Behavior & Organization* 126 (6): 58–71. DOI: 10.1016/j.jebo.2015.10.021. [11]

Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. 2019. "Food Deserts and the Causes of Nutritional Inequality." *Quarterly Journal of Economics* 134 (4): 1793–844. DOI: 10.1093/qje/qjz015. [18]

Antonakis, John, Anna T. Cianciolo, and Robert J. Sternberg. 2004. *The nature of leadership*. Sage Publications, Inc. [11]

Armand, Alex, Paul Atwell, and Joseph F. Gomes. 2020. "The Reach of Radio: Ending Civil Conflict through Rebel Demobilization." *American Economic Review* 110 (5): 1395–429. DOI: 10.1257/aer.20181135. [9, 10]

Bagwell, Kyle. 2007. "The Economic Analysis of Advertising." In *Handbook of Industrial Organization*. Edited by M. Armstrong and R. Porter. Vol. 3, Elsevier. Chapter 28, 1701–844. DOI: 10.1016/S1573-448X(06)03028-7. [10]

Bail, Christopher Andrew. 2016. "Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media." *Proceedings of the National Academy of Sciences* 113 (42): 11823–28. DOI: 10.1073/pnas.1607151113. [82]

Banerjee, Abhijit, Eliana La Ferrara, and Victor H. Orozco-Olvera. 2020. "The Entertaining Way to Behavioral Change: Fighting HIV with MTV." *Working paper*, [9, 10]

Bassi, Vittorio, and Imran Rasul. 2017. "Persuasion: A Case Study of Papal Influences on Fertility-Related Beliefs and Behavior." *American Economic Journal: Applied Economics* 9 (4): 250–302. DOI: 10.1257/app.20150540. [11]

Belk, Russell W., and Richard W. Pollay. 1985. "Images of Ourselves: The Good Life in Twentieth Century Advertising." *Journal of Consumer Research* 11 (4): 887. DOI: 10.1086/209024. [10]

Bénabou, Roland, Armin Falk, and Jean Tirole. 2020. "Narratives, Imperatives, and Moral Persuasion." *Working Paper*, URL: https://scholar.princeton.edu/sites/default/files/rbenabou/files/morals_september_15.pdf. [11]

Berg, Gunhild, and Bilal Zia. 2017. "Harnessing emotional connections to improve financial decisions: Evaluating the impact of financial education in mainstream media." *Journal of the European Economic Association* 15 (5): 1025–55. DOI: 10.1093/jeea/jvw021. [9]

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2018. "Behavioral Household Finance." In *Handbook of Behavioral Economics*. Vol. 1, Elsevier B.V., 177–276. DOI: 10.1016/bs.hesbe.2018.07.004. [9]

Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Katherine L. Milkman. 2015. "The Effect of Providing Peer Information on Retirement Savings Decisions." *Journal of Finance* 70 (3): 1161–201. [9]

Besley, Timothy, Thiemo Fetzer, and Hannes Mueller. 2019. "Terror and Tourism: The Economic Consequences of Media Coverage." *Working Paper*, 1–89. [24]

Bjorvatn, Kjetil, Alexander W. Cappelen, Linda Helgesson Sekei, Erik Ø. Sørensen, and Bertil Tungodden. 2020. "Teaching Through Television: Experimental Evidence on Entrepreneurship Education in Tanzania." *Management Science* 66 (6): 2308–25. DOI: 10.1287/mnsc.2019.3321. [9]

Blake, Thomas, Chris Nosko, and Steven Tadelis. 2015. "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment." *Econometrica* 83 (1): 155–74. DOI: 10.3982/ECTA12423. [10]

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." In *Journal of Machine Learning Research 3*. Vol. 3, Elsevier, 993–1022. [13]

Blouin, Arthur, and Sharun W. Mukand. 2019. "Erasing Ethnicity? Propaganda, Nation Building, and Identity in Rwanda." *Journal of Political Economy* 127 (3): 1008–62. DOI: 10.1086/701441. [16]

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2021. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Work in Progress*, 1–48. [9, 18, 22, 35, 49, 60]

Breiman, Leo. 2001. "Random Forest." *Machine Learning* 45 (10): 5–32. DOI: https://doi.org/10.1023/A:1010933404324. [7, 23]

Bronnenberg, Bart J., Jean-Pierre H Dubé, and Matthew Gentzkow. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review* 102 (6): 2472–508. DOI: 10.1257/aer.102.6.2472. [18]

Brown, Alexander L., and Joanna N. Lahey. 2015. "Small Victories: Creating Intrinsic Motivation in Task Completion and Debt Repayment." *Journal of Marketing Research* 52 (6): 768–83. DOI: 10.1509/jmr.14.0281. [13]

Bursztyn, Leonardo, and Davide Cantoni. 2016. "A Tear in the Iron Curtain: The Impact of Western Television on Consumption Behavior." *Review of Economics and Statistics* 98 (1): 25–41. DOI: 10.1162/REST. [10]

Bursztyn, Leonardo, Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao. 2018. "Status Goods: Experimental Evidence from Platinum Credit Cards*." *Quarterly Journal of Economics* 133 (3): 1561–95. DOI: 10.1093/qje/qjx048. [9]

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." *SSRN Electronic Journal*, DOI: 10.2139/ssrn.3580487. [11]

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2020. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 1–45. DOI: 10.1016/j.jeconom.2020.12.001. [9, 35]

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions*." *Quarterly Journal of Economics* 124 (4): 1639–74. [9]

Cavallo, Alberto. 2017. "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers." *American Economic Review* 107 (1): 283–303. DOI: 10.1257/aer.20160542. [16]

Cavell, Garrison C., Thomas G. Osenkowsky, David H. Layer, Skip Pizzi, and William T. Hayes. 2017. *National Association of Broadcasters Engineering Handbook*. Routledge, 1988. DOI: 10.4324/9781315680149. [16]

Chaisemartin, Clément de, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9): 2964–96. DOI: 10.1257/aer.20181169. [9, 35]

Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov. 2009. "Conspicuous Consumption and Race." *Quarterly Journal of Economics* 124 (2): 425–67. DOI: 10.1162/qjec.2009.124.2.425. [9]

Chen, Yuxin, Yogesh V. Joshi, Jagmohan S. Raju, and Z. John Zhang. 2009. "A Theory of Combative Advertising." *Marketing Science* 28 (1): 1–19. DOI: 10.1287/mksc.1080.0385. [11]

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. " Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark *." *Quarterly Journal of Economics* 129 (3): 1141–219. [9]

Chopra, Felix. 2021. "Media Persuasion and Consumption: Evidence from the Dave Ramsey Show." *SSRN Working Paper 3992358*, 115. DOI: http://dx.doi.org/10.2139/ssrn.3992358.

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2021. "Do People Demand Fact-Checked News? Evidence From U.S. Democrats." *ECONtribute Discussion Paper No. 121*, [37]

Coville, Aidan, Vincenzo Di Maro, Felipe Alexander Dunsch, and Siegfried Zottel. 2019. "The Nollywood Nudge: An Entertaining Approach to Saving." *Policy Research Working Paper 8920*, DOI: 10.1596/1813-9450-8920. [9]

Dahl, Gordon, and Stefano DellaVigna. 2009. "Does Movie Violence Increase Violent Crime?" *Quarterly Journal of Economics* 124 (2): 677–734. DOI: 10.1162/qjec.2009.124.2.677. [10]

Davies, Emma, and Stephen E.G. Lea. 1995. "Student attitudes to student debt." *Journal of Economic Psychology* 16 (4): 663–79. DOI: 10.1016/0167-4870(96)80014-6. [8, 39]

Della Vigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova, and Ekaterina Zhuravskaya. 2014. "Cross-border media and nationalism: Evidence from serbian radio in croatia." *American Economic Journal: Applied Economics* 6 (3): 103–32. DOI: 10.1257/app.6.3.103. [10]

Della Vigna, Stefano, and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *Quarterly Journal of Economics* 3 (122): 1187–234. DOI: 10.1162/qjec.122.3.1187. [10, 29]

DellaVigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics* 2 (1): 643–69. DOI: 10.1146/annurev.economics.102308.124309. [10]

DellaVigna, Stefano, and Eliana La Ferrara. 2015. "Economic and Social Impacts of the Media." In *Handbook of Media Economics*. Edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg. Vol. 1, Elsevier B.V. Chapter 19, 723–68. DOI: http://dx.doi.org/10.1016/B978-0-444-63685-0.00019-X. [29]

Dori-Hacohen, Gonen. 2019. ""I can do math, but I'm not that smart. I'm not brilliant."" In *The Construction of 'Ordinariness' across Media Genres*. Edited by Anita Fetzer and Elda Weizman, 133–55. DOI: 10.1075/pbns.307.06dor. [12]

Drexler, Alejandro, Greg Fischer, and Antoinette Schoar. 2014. "Keeping It Simple: Financial Literacy and Rules of Thumb." *American Economic Journal: Applied Economics* 6 (2): 1–31. DOI: 10.1257/app.6.2.1. [13]

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2018. "Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession." *Marketing Science* 37 (1): 22–53. DOI: 10.1287/mksc.2017.1047. [17]

Durante, Ruben, and Brian Knight. 2012. "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy." *Journal of the European Economic Association* 10 (3): 451–81. DOI: 10.1111/j.1542-4774.2011.01060.x. [10]

Durante, Ruben, Paolo Pinotti, and Andrea Tesei. 2019. "The Political Legacy of Entertainment TV." *American Economic Review* 109 (7): 2497–530. DOI: 10.1257/aer.20150958. [10, 16]

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2010. "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?" *Quantitative Marketing and Economics* 8 (2): 207–39. DOI: 10.1007/s11129-009-9073-0. [17]

Eliaz, Kfir, and Ran Spiegler. 2020. "A Model of Competing Narratives." *American Economic Review* 110 (12): 3786–816. DOI: 10.1257/aer.20191099. [11]

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–85. DOI: 10.1257/aer.101.7.3253. [10]

Falk, Armin, and Markus Knell. 2004. "Choosing the Joneses: Endogenous Goals and Reference Standards." *Scandinavian Journal of Economics* 106 (3): 417–35. DOI: 10.1111/j.0347-0520.2004.00370.x. [9]

Fernandes, Daniel, John G. Lynch, and Richard G. Netemeyer. 2014. "Financial Literacy, Financial Education, and Downstream Financial Behaviors." *Management Science* 60 (8): 1861–83. DOI: 10.1287/mnsc.2013.1849. [9]

Frank, Robert H. 1985. *Choosing the Right Pond: Human Behavior and the Quest for Status*. Oxford University Press. [9]

Gentzkow, Matthew. 2006. "Television and Voter Turnout*." *Quarterly Journal of Economics* 121 (3): 931–72. DOI: 10.1162/qjec.121.3.931. [10]

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as Data." *Journal of Economic Literature* 57 (3): 535–74. DOI: 10.1257/jel.20181020. [13, 82]

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai. 2021. "Household Finance." *Journal of Economic Literature* 59 (3): 919–1000. DOI: 10.1257/jel.20201461. [6]

Goodman-Bacon, Andrew. 2019. "Difference-in-Differences With Variation in Treatment Timing." *Working paper*, [9, 18, 35]

Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo. 2009. "Consumer Shopping Behavior: How Much Do Consumers Save?" *Journal of Economic Perspectives* 23 (2): 99–120. DOI: 10.1257/jep.23.2.99. [31]

Haaland, Ingar, and Christopher Roth. 2021. "Beliefs about Racial Discrimination and Support for Pro-Black Policies." *Review of Economics and Statistics*, (03): 1–38. DOI: 10.1162/rest_a_01036. [37]

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. Forthcoming. "Designing Information Provision Experiments." *Journal of Economic Literature*, [9]

Hastings, Justine S., Brigitte C. Madrian, and William L. Skimmyhorn. 2013. "Financial Literacy, Financial Education, and Economic Outcomes." *Annual Review of Economics* 5 (1): 347–73. DOI: 10.1146/annurev-economics-082312-125807. [9, 38]

Heffetz, Ori. 2011. "A Test of Conspicuous Consumption: Visibility and Income Elasticities." *Review of Economics and Statistics* 93 (4): 1101–17. DOI: 10.1162/REST{_}a{_}00116. [9]

Hendricks, John Allen, and Bruce Mims. 2018. *The Radio Station: Broadcasting, Podcasting, and Streaming*. 10th edition. Routledge, 482. [19]

Hjarvard, Stig. 2008. "The Mediatization of Society." *Nordicom Review* 29 (2): 102–31. DOI: 10.1515/nor-2017-0181. [10]

Hjarvard, Stig. 2013. *The Mediatization of Culture and Society*. Routledge. [10]

Imai, Kosuke, and In Song Kim. 2019. "When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data?" *American Journal of Political Science* 63 (2): 467–90. DOI: 10.1111/ajps.12417. [36]

Jensen, Robert, and Emily Oster. 2009. "The Power of TV: Cable Television and Women's Status in India *." *Quarterly Journal of Economics* 124 (3): 1057–94. DOI: 10.1162/qjec.2009.124.3.1057. [10]

Jones, B. F., and B. A. Olken. 2005. "Do Leaders Matter? National Leadership and Growth Since World War II." *Quarterly Journal of Economics* 120 (3): 835–64. DOI: 10.1093/qje/120.3.835. [11]

Karlan, Dean, Aishwarya Lakshmi Ratan, and Jonathan Zinman. 2014. "Savings by and for the Poor: A Research Review and Agenda." *Review of Income and Wealth* 60 (1): 36–78. DOI: 10.1111/roiw.12101. [6]

Kasampalis, S, P I Lazaridis, Z D Zaharis, A Bizopoulos, S Zettas, and J Cosmas. 2013. "Comparison of Longley-Rice, ITM and ITWOM propagation models for DTV and FM broadcasting." In *2013 16th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, 1–6. [16]

Kerkhof, Anna. 2020. "Advertising and Content Differentiation: Evidence from YouTube." *CESifo Working Paper No. 8697*, [13]

Kettle, Keri L., Remi Trudel, Simon J. Blanchard, and Gerald Häubl. 2016. "Repayment concentration and consumer motivation to get out of debt." *Journal of Consumer Research* 43 (3): 460–77. DOI: 10.1093/jcr/ucw037. [13]

La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea. 2012. "Soap Operas and Fertility: Evidence from Brazil." *American Economic Journal: Applied Economics* 4 (4): 1–31. DOI: 10.1257/app.4.4.1. [9, 10]

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–78. DOI: 10.1162/003355397555253. [6]

Le Rest, Kévin, David Pinaud, Pascal Monestiez, Joël Chadoeuf, and Vincent Bretagnolle. 2014. "Spatial leave-one-out cross-validation for variable selection in the presence of spatial autocorrelation." *Global Ecology and Biogeography* 23 (7): 811–20. DOI: 10.1111/geb.12161. [23]

Lewis, Randall, and Dan Nguyen. 2015. "Display advertising's competitive spillovers to consumer search." *Quantitative Marketing and Economics* 13 (2): 93–115. DOI: 10.1007/s11129-015-9155-0. [10]

Lodish, Leonard M., Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, and Mary Ellen Stevens. 1995. "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments." *Journal of Marketing Research* 32 (2): 125. DOI: 10.2307/3152042. [10]

Lusardi, Annamaria, and Olivia S. Mitchell. 2007. "Baby Boomer retirement security: The roles of planning, financial literacy, and housing wealth." *Journal of Monetary Economics* 54 (1): 205–24. DOI: 10.1016/j.jmoneco.2006.12.001. [9]

Lusardi, Annamaria, Olivia S. Mitchell, and Noemi Oggero. 2020. "Debt and Financial Vulnerability on the Verge of Retirement." *Journal of Money, Credit and Banking* 52 (5): 1005–34. DOI: 10.1111/jmcb.12671. [6]

Lusk, Jayson L, and Kathleen Brooks. 2011. "Who Participates in Household Scanning Panels?" *American Journal of Agricultural Economics* 93 (1): 226–40. DOI: 10.1093/ajae/aaq123. [17]

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior*." *Quarterly Journal of Economics* 116 (4): 1149–87. [9]

Marshall, Alfred. 1919. *Industry and Trade: A Study of Industrial Technique and Business Organization; and of Their Influences on the Conditions of Various Classes and Nations*. MacMillan, and Co., London. [10]

Martin, Gregory J., and Ali Yurukoglu. 2017. "Bias in cable news: Persuasion and polarization." *American Economic Review* 107 (9): 2565–99. DOI: 10.1257/aer.20160812. [29]

Mian, Atif, Amir Sufi, and Emil Verner. 2017. "Household Debt and Business Cycles Worldwide*." *Quarterly Journal of Economics* 132 (4): 1755–817. DOI: 10.1093/qje/qjx017. [6]

MIT Election Data and Science Lab. 2018. *County Presidential Election Returns 2000-2020*. Version V9. Harvard Dataverse. DOI: 10.7910/DVN/VOQCHQ. [18, 36, 52]

Müller, Karsten, and Carlo Schwarz. 2020. "From Hashtag to Hate Crime: Twitter and Anti-Minority Sentiment." *SSRN Electronic Journal*, [11]

Okuyama, Yoko. 2019. "Toward Better Informed Decision-Making: The Impacts of a Mass Media Campaign on Women' s Outcomes in Occupied Japan." *Working Paper*, 1–57. [10]

Olken, Benjamin A. 2009. "Do Television and Radio Destroy Social Capital? Evidence from Indonesian Villages." *American Economic Journal: Applied Economics* 1 (4): 1–33. DOI: 10.1257/app.1.4.1. [7, 9, 16]

Perebinossoff, Philippe, Brian Gross, and Lynne S. Gross. 2005. *Programming for TV, Radio, and the Internet*. Routledge, 345. DOI: 10.4324/9780080480466. [14]

Richins, Marsha L. 1987. "Media, Materialism, and Human Happiness." *ACR North American Advances*, [10]

Richins, Marsha L., and Scott Dawson. 1992. "A Consumer Values Orientation for Materialism and Its Measurement: Scale Development and Validation." *Journal of Consumer Research* 19 (3): 303. DOI: 10.1086/209304. [8, 39]

Roberts, David R., Volker Bahn, Simone Ciuti, Mark S. Boyce, Jane Elith, Gurutzeta Guillera-Arroita, Severin Hauenstein, José J. Lahoz-Monfort, Boris Schröder, Wilfried Thuiller, David I. Warton, Brendan A. Wintle, Florian Hartig, and Carsten F. Dormann. 2017. "Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure." *Ecography* 40 (8): 913–29. DOI: 10.1111/ecog.02881. [23]

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using Models to Persuade." *American Economic Review* 111 (1): 276–323. DOI: 10.1257/aer.20191074. [11]

Shapiro, Bradley T., Günter J. Hitsch, and Anna E. Tuchman. 2021. "TV Advertising Effectiveness and Profitability: Generalizable Results From 288 Brands." *Econometrica* 89 (4): 1855–79. DOI: 10.3982/ECTA17674. [10]

Shiller, Robert J. 2020. *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton/Oxford: Princeton University Press. [11, 12]

Sweeting, Andrew. 2015. "Radio." *Handbook of Media Economics* 1: 341–96. DOI: 10.1016/B978-0-444-62721-6.00008-1. [13]

Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28 (2): 3–28. DOI: 10.1257/jep.28.2.3. [24]

Wager, Stefan, and Susan Athey. 2018. "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests." *Journal of the American Statistical Association* 113 (523): 1228–42. DOI: 10.1080/01621459.2017.1319839. [24]

Wang, Tianyi. 2021. "Media, Pulpit, and Populist Persuasion: Evidence from Father Coughlin." *American Economic Review* 111 (9): 3064–92. DOI: 10.1257/aer.20200513. [10, 11, 29]

Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics* 129 (4): 1947–94. DOI: 10.1093/qje/qju020. [10, 16, 29]

YouGov. 2021a. "YouGov America: The Clark Howard Show." https://today.yougov.com/topics/media/explore/radio_program/The_Dave_Ramsey_Show (accessed 10-19-2021). [11]

YouGov. 2021b. "YouGov America: The Dave Ramsey Show." https://today.yougov.com/topics/media/explore/radio_program/The_Clark_Howard_Show (accessed 10-19-2021). [11, 28]

Summary of the Appendices

Appendix Section 1.A contains details about the data sources and the construction of variables. Section 1.A.2 provides additional information on the data and procedures used to obtain the radio coverage indicator.

Appendix Section 1.B contains additional descriptive material. Figure 1.B.1 indicates the number of new affiliates by year. Figure 1.B.2 shows the Google Trend's popularity of the *Dave Ramsey Show* and *Sean Hannity* over time. Figure 1.B.3 presents the spatial distribution of affiliated radio stations across the US. Figure 1.B.4 and Figure 1.B.5 show DMA-level summary statistics for the number of Nielsen panelists and household expenditures.

Appendix Section 1.C contains additional robustness checks. Figure 1.C.1 presents event-study estimates using the imputation estimator proposed by Borusyak, Jaravel, and Spiess (2021).

Appendix Section 1.D contains additional analyses. Figure 1.D.1 presents the distribution of effects on household expenditures obtained from a repeated assignment of counterfactual market entry dates. Table 1.D.1 shows estimates of the effect of the radio show on political outcomes.

Appendix Section 1.E contains additional results related to the content analysis. Section 1.E.1 presents qualitative evidence on the radio show's narrative, advice and content. Section 1.E.2 presents quantitative evidence from text analysis of about 3,000 hours of content uploaded by the *Dave Ramsey Show* on its YouTube channel. Section 1.E.2.1 presents the YouTube data and how I prepare the text data for text analysis. Section 1.E.2.2 contains topic model estimates from Latent Dirichlet Allocation. Section 1.E.2.3 contains additional results on the most frequently spoken non-stopwords and keywords used to describe the videos (Table 1.E.1) and the top correlates of the word "debt" (Figure 1.E.2).

Appendix Section 1.F provides supplementary material for the experimental part of the paper (discussed in Section 1.6). I provide information about research transparency in Section 1.F.1, including a discussion of the preregistration, ethical approval, data and code availability, and a declaration of no conflict of interest. Section 1.F.2 contains additional figures and tables. In particular, a design overview (Figure 1.F.1), a comparison of sample characteristics to the general population (Table 1.F.1), a test of balance (Table 1.F.2), a test of balance for demographics elicited post-treatment (Table 1.F.3). The correlation between consumption and debt attitudes and self-reported behavior are shown in Table 1.F.4. Table 1.F.5 presents the main results without the inclusion of control variables. Table 1.F.6 presents treatment effects on individual items used to construct the consumption and debt attitude indices in the main experiment. Table 1.F.7 presents treatment effects on individual items used to construct the consumption and debt attitude indices in the follow-up survey. Table 1.F.8 tests for differential attrition across treatment arms between the main experiment and the follow-up survey. Table 1.F.9 presents a test of balance of covariates across treatment arms in the follow-up survey. Section 1.F.4 and Section 1.F.6 contain the original instructions used in the main experiment and the obfuscated follow-up survey, respectively.

Appendix 1.A Data

1.A.1 Data sources

Table 1.A.1. Data sources

Variables	Source	Comment
	Dependent variable	
Household expenditures, number of products purchased, other household-level outcomes based on UPC-level purchase records	Nielsen Homescan Data	Monthly household-level statistics result from aggregating purchase records across individual shopping trips
	Radio coverage	
Signal strength, free-space signal strength	Own calculations	Derived from an implementation of the Longley-Rice/Irregular Terrain Model
Radio coverage	Ramsey Media, own calculations	Construction as described in Section 1.3, combining signal strength measures and information about the timing of market entry. This variable varies at the zip code-month level.
	Control variables	
Household-level covariates	Nielsen Homescan Data	Self-reported sociodemographic variables, elicited each fall
Unemployment rate	US Bureau of Labor Statistics	The unemployment rate varies at the county-month level
Urbanization	US Census Bureau	Share of the zip code population living in urban areas. Based on data from the H002 Urban and Rural Summary File 1.
House prices	Zillow Group	This is the Zillow Home Price Index. Data series are obtained at the zip code and the county-month level. Available at: https://www.zillow.com/research/data/

Variables	Source	Comment
Christian share	US Religion Census	US Religion Census: Religious Congregations and Membership Study, 2010 (County File), accessed: October 2019.
County-level demographics	US Census (2000, 2010), American Community Survey	Vary at the county-year level
Radio transmitter characteristics		
Transmitter height, power, frequency, and location	Federal Communications Commission (FCC)	Power in kilowatt, height in meter, frequency in MHz and coordinates of the transmitter in NAD83 coordinates. Data obtained using the AM and FM Query tools available at: https://www.fcc.gov/licensing-databases/search-fcc-databases; accessed February 2019
Geographical variables		
State, county and zip code boundaries	US Census Bureau	Shapefiles for state, county and ZCTA representation of 5-digit zip codes (1:500k) in WGS84 coordinates. Data available at: https://www2.census.gov/geo/tiger/GENZ2017/shp/
Boundaries for Designated Market Areas (DMAs)	Nielsen	Based on a cross-walk from Designated Market Areas to US counties available from Nielsen.
Latitude and longitude of the geographic center of administrative units		Derived from the corresponding shapefiles using the Python package geopandas after applying a distance-preserving projection
Terrain elevation	Global Land One-km Base Elevation Project (GLOBE)	Height above mean sea levels (in meters). Available at: https://www.ngdc.noaa.gov/mgg/topo/globe.html; accessed October 2020
Other variables		
Political outcomes (turnout, vote shares)	MIT Election Data and Science Lab (2018)	County-level electoral results for the Presidential elections between 2000–2016

1.A.2 Radio coverage

This section provides additional details on how I determined the spatial radio coverage of affiliated radio stations.

The information on the affiliated radio stations of the *Dave Ramsey Show* included their current their call sign, frequency, and the DMA, state and city where the radio station is located. However, radio stations often change their call sign when they switch to a new format. To obtain time-invariant identifiers, I manually match all affiliated radio stations with the FCC transmitter identifier of their primary transmitter ("Facility ID"). Moreover, many radio stations operate multiple transmitters in different locations to increase their service area and provide better radio coverage. For all affiliated stations, I thus obtain a complete list of their secondary transmitters from the FCC, including the exact date when the secondary transmitter started to broadcast. In my analysis, I include the radio coverage of secondary transmitters after the latter of (i) the date when their primary transmitter started to broadcast the radio show and (ii) the date the secondary transmitter actually started to broadcast.

For each transmitter, I then obtain the geographic coordinates of their location and the technical parameters needed for the signal propagation models. In the case of the Longley-Rice/Irregular Terrain Model, these parameters include the effectively radiated power (in kilowatts), the height of the transmitter antenna above ground levels (in meters), and the broadcast frequency (in MHz). The model also requires topographic information on the elevation profile to account for the effect of obstructions that block line-of-sight transmission. I use data from the Global 30 Arc-Second Elevation Database.

I then use the Longley-Rice/Irregular Terrain Model to calculate the path loss (in dB) between pairs of receiver and transmitter locations. The program code was obtained from Benjamin Olken. As the residency of Nielsen households is known up to the 5-digit zip code, I use the geographic coordinates of the centroid of zip codes as potential receiver locations. For each transmitter, I calculate the signal loss for all zip codes within 600 km of the transmitter's location. In addition, I calculate the free-space path loss using the same parameters. I then deduct the path loss from the transmitter signal strength to obtain the receiver signal strength. Whenever a zip code receives a radio signal from multiple transmitters, I follow the literature and use the maximum receiver signal strength.

For county-level analyses, I calculate the share of the population with access to the *Dave Ramsey Show*. Specifically, I use a signal strength threshold of 50 dBμV/m to classify zip codes as having radio coverage. I calculate the share of the county population accounted for by zip codes with radio coverage.

Figure 1.A.1 provides an example of the zip code-level variation in radio signal strength. The figure plots radio signal strength (in deciles) in California at the end of 2012.

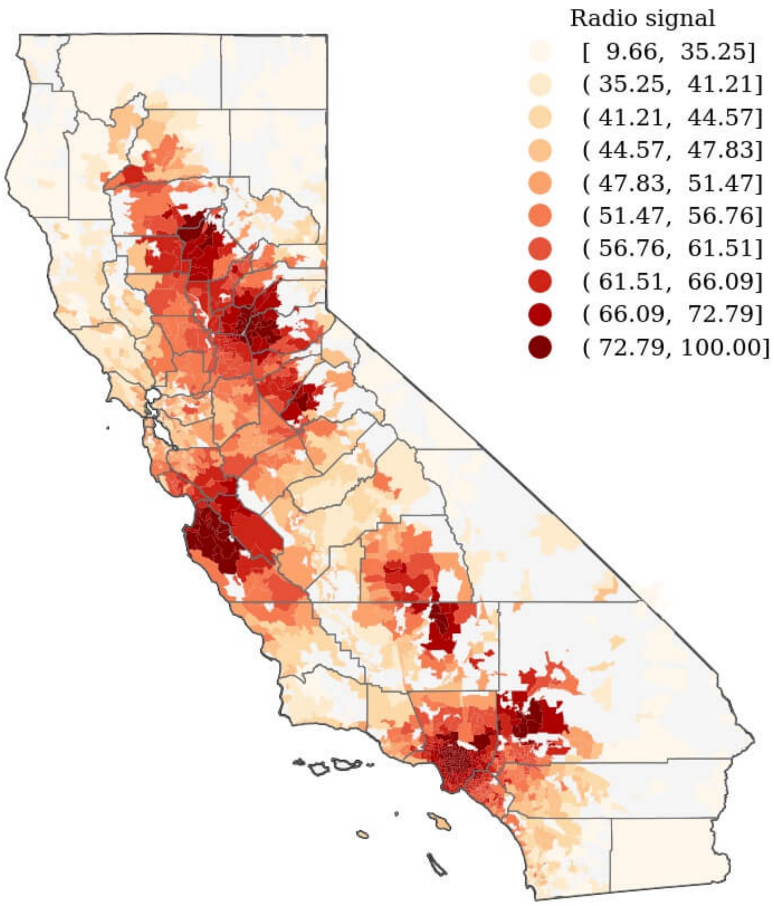

Figure 1.A.1. Radio coverage at the zip code level: Example

Notes: This figure displays the radio signal strength (in dBμV/m) across zip codes in California as of 2012. The radio signal is the maximum signal strength across all transmitters of affiliated radio stations and capped at 100 dBμV/m.

Appendix 1.B Descriptives

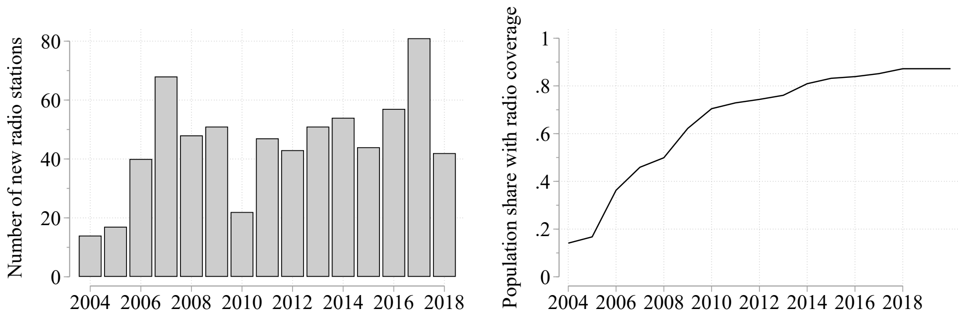

Figure 1.B.1. Expansion of the affiliate network over time

Notes: This panel on the left displays the number of new affiliated radio stations starting to broadcast the *Dave Ramsey Show* over time. The panel on the right plots the share of the US population with radio coverage from an affiliated radio stations over time.

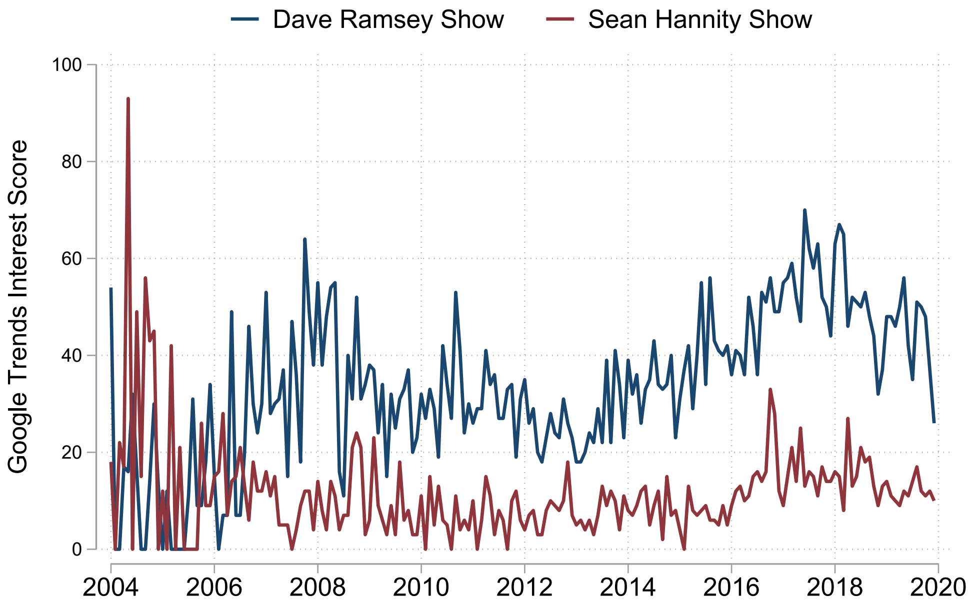

Figure 1.B.2. Popularity of the Dave Ramsey Show as measured by Google searches

Notes: This figure uses monthly Google Trends data for the period from January 1, 2004, to December 31, 2019. For each month, the figure indicates the interested in the two topics "The Dave Ramsey Show" and "The Sean Hannity Show" as determined by Google searches related to these topics. The Google Trends data is normalized to a scale ranging from 0 to 100, where larger values indicate more searches. The data was obtained on June 17, 2021, from https://trends.google.com.

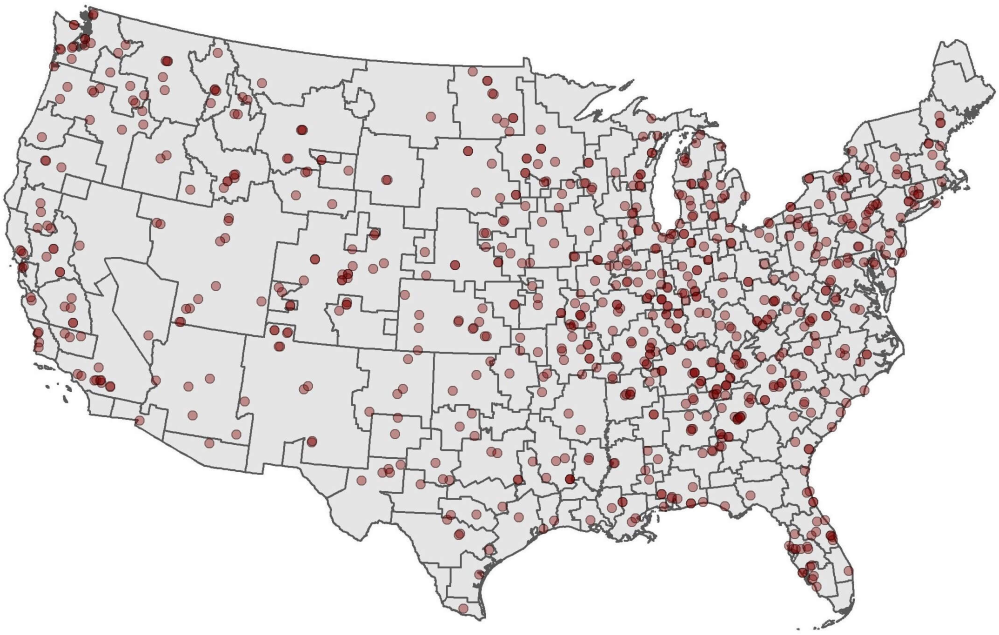

Figure 1.B.3. Transmitter locations of affiliated radio stations

Notes: This map plots the locations of the transmitters of all radio stations broadcasting the radio show together with the boundaries of Nielsen's Designated Market Areas (DMAs).

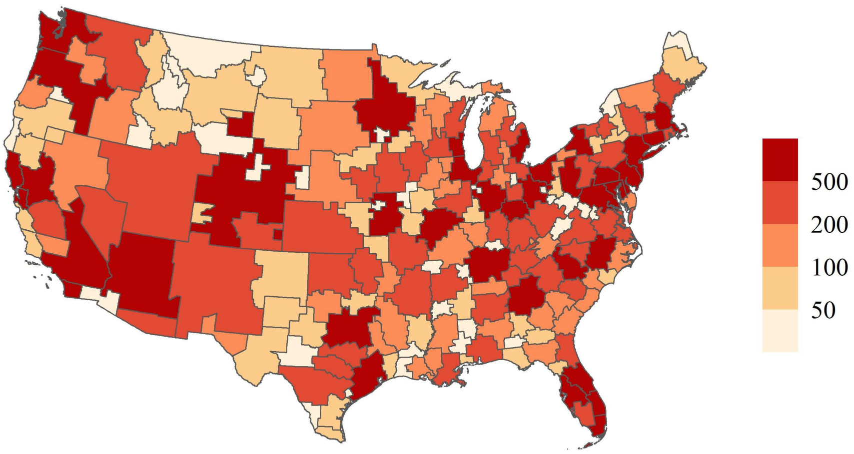

Figure 1.B.4. Nielsen panelists by Designated Market Area

Notes: This map shows the total number of Nielsen panelists in 2017 by Designated Market Area.

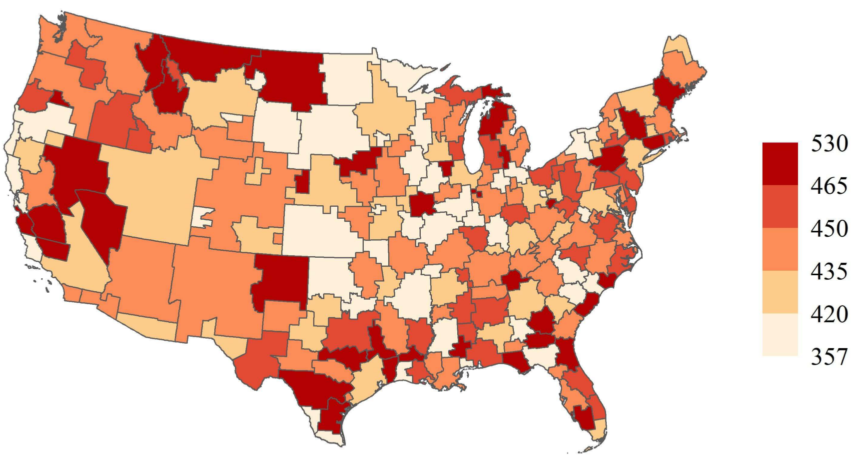

Figure 1.B.5. Monthly expenditures by Designated Market Area

Notes: This map shows the average monthly expenditure of Nielsen panelists (in $) in 2017 by Designated Market Area.

Appendix 1.C Robustness checks

Table 1.C.1. Poisson regression – Number of purchased items

	Dependent variable: Number of purchased products			
	(1)	(2)	(3)	(4)
Radio show	-0.0187***	-0.0182***	-0.0224***	-0.0240***
	(0.0031)	(0.0031)	(0.0031)	(0.0035)
N	3,744,054	3,744,054	3,407,688	3,407,688
Pseudo R^2	0.517	0.520	0.521	0.523
Mean of dep. var.	83.30	83.30	83.06	83.06
Household & Time FEs	Yes	Yes	Yes	Yes
Household controls		Yes	Yes	Yes
Local economic conditions			Yes	Yes
State x Time FEs				Yes

Notes: This table show Poisson regression estimates using 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable is the number of purchased items per month. "Radio show" is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. Individual controls include the log of household income, age indicators, household size, married indicator and employment status indicators (full-time, part-time, unemployed). Local economic conditions comprise controls for house prices and the unemployment rate. Robust standard errors clustered by zip code are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

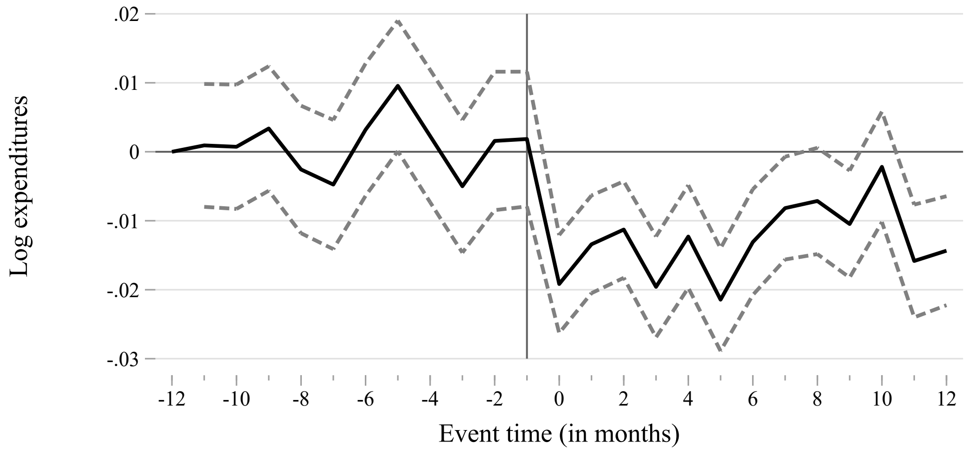

Figure 1.C.1. Robustness to treatment effect heterogeneity: Borusyak, Jaravel, and Spiess (2021) imputation estimator

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable are log expenditures. The omitted category is 12 months before market entry. Estimates of the treatment effect dynamics are obtained from the imputation estimator proposed by Borusyak, Jaravel, and Spiess (2021). The estimator includes household and year-month fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level.

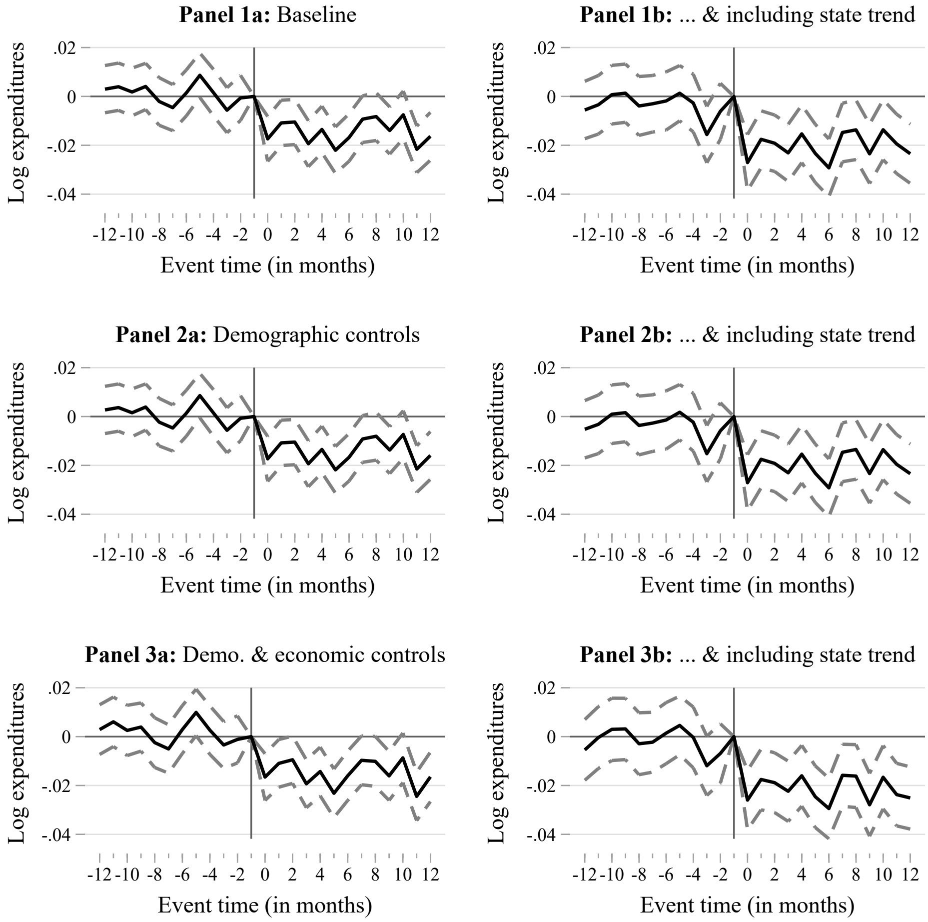

Figure 1.C.2. Robustness: Control variables and state-specific time trends

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable are log expenditures. The baseline specification in Panel 1a includes event time indicators, household fixed effects and year-month fixed effects. The month before market entry serves as the omitted category. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panel 1b, 2b and 3b include state×year-month fixed effects to the specification in Panel 1a, 2a and 3a, respectively. Panel 2a includes time-varying household-level demographic controls. Panel 3a includes time-varying household-level demographic controls and proxies for local economic conditions, including monthly house prices (zip code) and the monthly unemployment rate (county level).

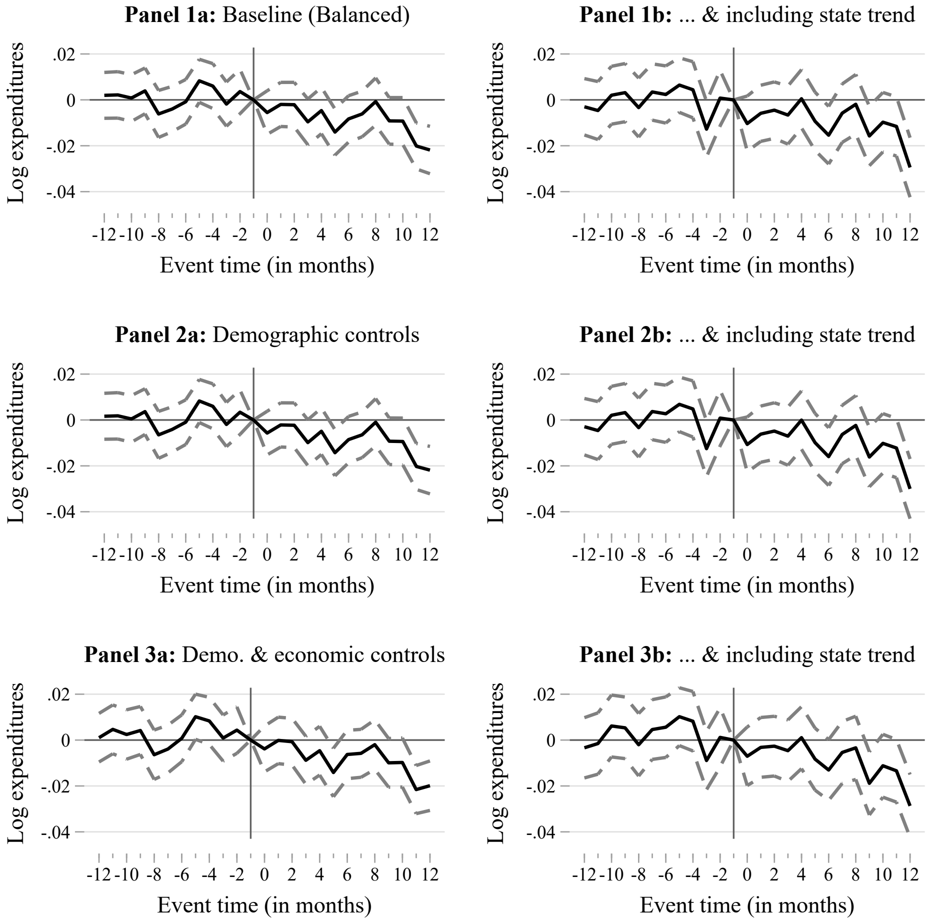

Figure 1.C.3. Robustness: Balanced sample

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. All panels use a balanced sample in event time. The dependent variable are log expenditures. The baseline specification in Panel 1a includes event time indicators, household fixed effects and year-month fixed effects. The month before market entry serves as the omitted category. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panel 1b, 2b and 3b include state×year-month fixed effects to the specification in Panel 1a, 2a and 3a, respectively. Panel 2a includes time-varying household-level demographic controls. Panel 3a includes time-varying household-level demographic controls and proxies for local economic conditions, including monthly house prices (zip code) and the monthly unemployment rate (county level).

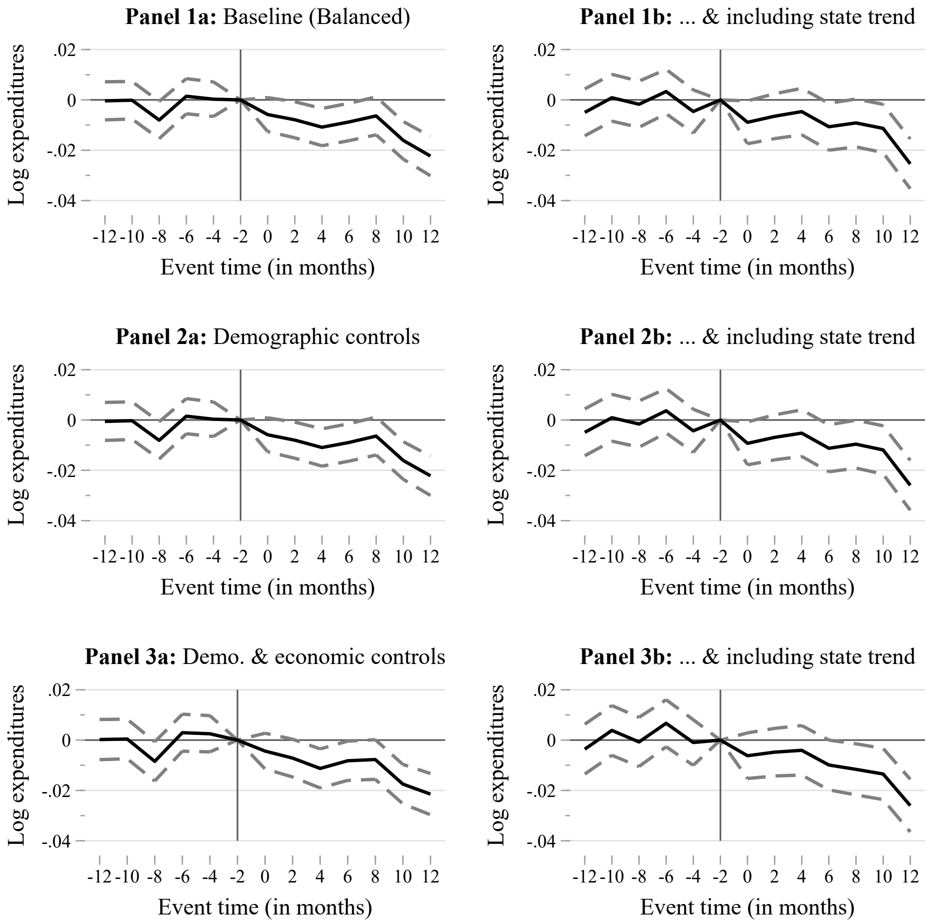

Figure 1.C.4. Robustness: Balanced sample with binned event time indicators

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. All panels use a balanced sample in event time. The dependent variable are log expenditures. The baseline specification in Panel 1a includes event time indicators, household fixed effects and year-month fixed effects. The month before market entry serves as the omitted category. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panel 1b, 2b and 3b include state×year-month fixed effects to the specification in Panel 1a, 2a and 3a, respectively. Panel 2a includes time-varying household-level demographic controls. Panel 3a includes time-varying household-level demographic controls and proxies for local economic conditions, including monthly house prices (zip code) and the monthly unemployment rate (county level).

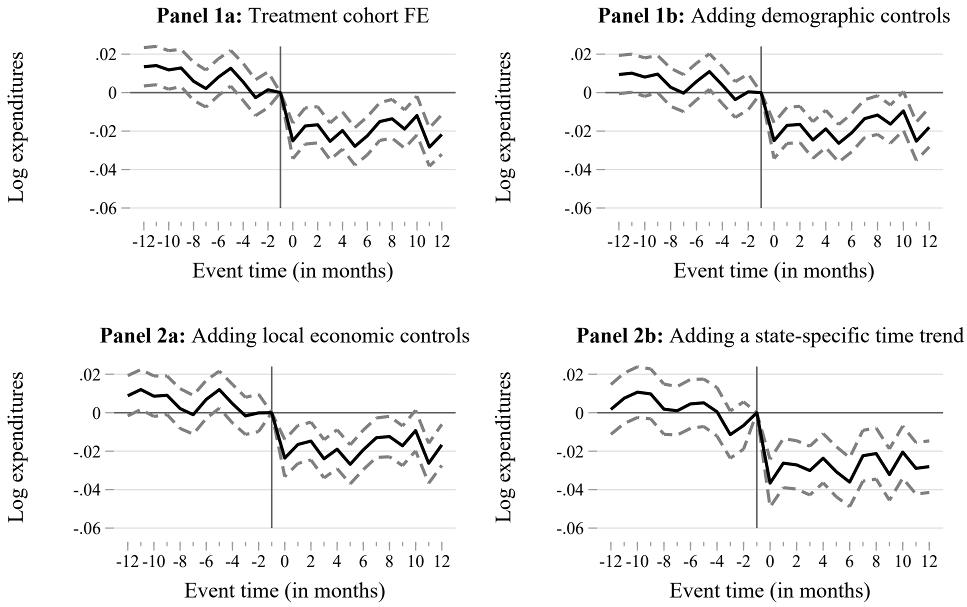

Figure 1.C.5. Robustness: Treatment cohort instead of household fixed effects

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable are log expenditures. All regressions include event time indicators and year-month fixed effects. Moreover, all regressions include treatment cohort fixed effects (defined by the year-month a household is first treated) and zip code fixed effects instead of household fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panel 1b adds time-varying household-level demographic controls to the specification from Panel 1a. Panel 2a further adds proxies for local economic (house prices, unemployment rate) to the set of control variables. Panel 2b includes the full set of controls and state×year-month fixed effects.

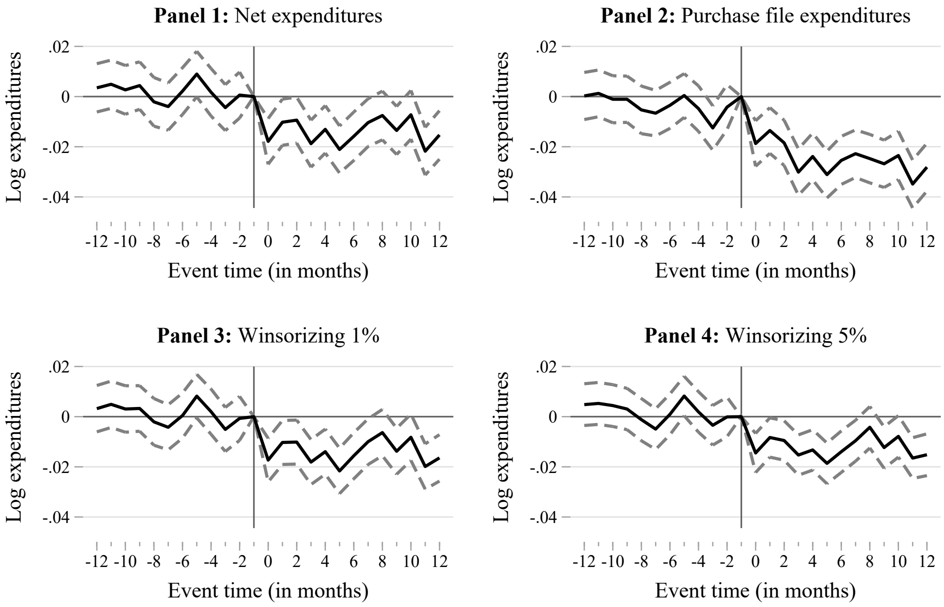

Figure 1.C.6. Event study: Alternative measures of household expenditures

Notes: This figure uses 2004–2019 Nielsen Homescan data at the household-by-month level. All regressions include event time indicators, household and year-month fixed effects, the full set of controls, and state×year-month fixed effects. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level. Panels differ in how monthly household expenditures are constructed. Panel 1 uses monthly expenditures net of the value of redeemed coupons as the dependent variable. Panel 2 uses the sum of all expenditures recorded in the Nielsen Homescan purchase files, excluding data supplied to Nielsen from retailers. Panel 3 and 4 winsorize household expenditures at the 1% and 5% level, respectively.

Table 1.C.2. Robustness to using alternative measures of household expenditures

	(1) Net expenditures	(2) Purchase file expenditures	(3) Winsorizing 1%	(4) Winsorizing 5%
Radio coverage	-0.016***	-0.032***	-0.016***	-0.015***
	(0.003)	(0.003)	(0.003)	(0.002)
N	3,399,591	3,399,566	3,407,700	3,407,700
R^2	0.527	0.551	0.537	0.549
Mean of dep. var.	6.169	5.639	6.190	6.201
Household & Time FEs	Yes	Yes	Yes	Yes
Household controls	Yes	Yes	Yes	Yes
Local economic conditions	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable is the log of household expenditures, where expenditures are constructed as indicated by the column header. Specifically, column 1 uses monthly expenditures net of the value of redeemed coupons. Column 2 uses the sum of all expenditures recorded in the Nielsen Homescan purchase files, excluding data supplied to Nielsen from retailers. Columns 3 and 4 winsorize the household expenditures at the 1% and 5% level, respectively. Robust standard errors clustered at the zip code level are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

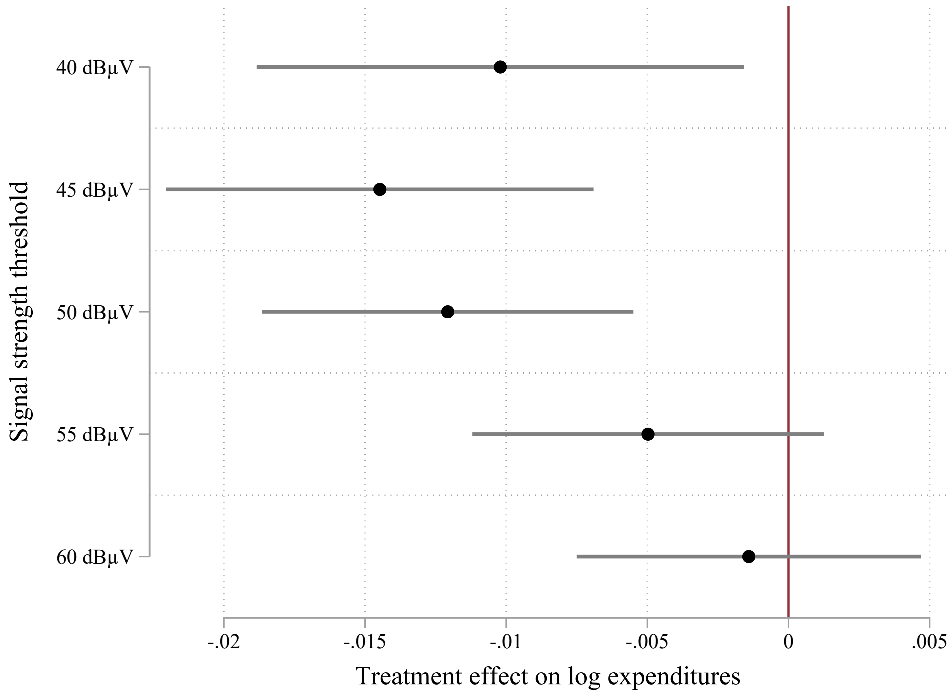

Figure 1.C.7. Robustness – Alternative signal strength thresholds

Notes: This figure plots estimates of the baseline model (equation 1.4.1) using alternative thresholds to binarize the continuous signal strength measure. The dependent variable are log household expenditures. All regressions include household and year-month fixed effects, state×year-month fixed effects and the set of time-varying controls. 95% confidence intervals are constructed from robust standard errors clustered at the zip code level.

Table 1.C.3. Robustness – Excluding households based on when they receive radio coverage

	Excluded treatment cohorts:							
	(1) 04/05	(2) 06/07	(3) 08/09	(4) 10/11	(5) 12/13	(6) 14/15	(7) 16/17	(8) 18/19
Radio coverage	-0.013*** (0.003)	-0.015*** (0.003)	-0.012*** (0.003)	-0.012*** (0.003)	-0.011*** (0.003)	-0.012*** (0.003)	-0.012*** (0.003)	-0.013*** (0.003)
N	3,744,066	3,020,964	3,462,420	3,583,859	3,641,719	3,511,737	3,720,165	3,696,707
R^2	0.518	0.520	0.518	0.518	0.518	0.517	0.518	0.517
Household FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in all regressions are log household expenditures. Robust standard errors clustered at the zip code level are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.4. Robustness: Log expenditures – Varying the sample of DMAs

	Dependent variable: log (Expenditures)						
	Excluded DMA ranks				Included DMA ranks		
	(1) 1–50	(2) 51–100	(3) 101–150	(4) 150–210	(5) 1–50	(6) 51–100	(7) 101–210
Radio coverage	-0.010**	-0.012***	-0.015***	-0.013***	-0.014***	-0.015**	-0.002
	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)	(0.006)	(0.008)
N	1,209,747	3,006,281	3,447,726	3,568,444	2,534,319	737,785	471,962
R^2	0.521	0.517	0.518	0.518	0.517	0.523	0.517
Household FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in all regressions are log household expenditures. Robust standard errors clustered at the zip code level and shown in parentheses. Nielsen DMA market rankings are from 2017.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.5. Robustness: Log items – Varying the sample of DMAs

	Dependent variable: log (Number of purchased items)						
	Excluded DMA ranks				Included DMA ranks		
	(1) 1–50	(2) 51–100	(3) 101–150	(4) 150–210	(5) 1–50	(6) 51–100	(7) 101–210
Radio coverage	-0.024***	-0.013***	-0.017***	-0.017***	-0.014***	-0.030***	-0.014*
	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)	(0.007)	(0.008)
N	1,206,284	2,998,991	3,439,762	3,559,606	2,528,597	735,890	470,394
R^2	0.538	0.540	0.542	0.542	0.541	0.542	0.532
Household FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in all regressions is the log of the number of purchased products per month. Robust standard errors clustered at the zip code level and shown in parentheses. Nielsen DMA market rankings are from 2017.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.6. Robustness: Expenditures – Excluding counties with affiliates and areas close to Nashville

	Dependent variable: log (Expenditures)					
	Drop zip codes close to Nashville		Drop counties with affiliate stations		Apply both restrictions	
	(1)	(2)	(3)	(4)	(5)	(6)
Radio coverage	-0.013***	-0.012***	-0.013***	-0.009**	-0.013***	-0.011**
	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.005)
N	3,345,355	3,048,109	2,314,720	2,036,495	2,050,384	1,804,011
R^2	0.519	0.525	0.520	0.527	0.521	0.529
Mean of dep. var.	6.190	6.191	6.186	6.187	6.191	6.192
Household & Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
Household controls		Yes		Yes		Yes
Local economic conditions		Yes		Yes		Yes
State x Time FEs		Yes		Yes		Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in all regressions are log household expenditures. Columns 1–2 exclude households residing in zip codes within 500 km of Nashville, Tennessee. Columns 3–4 exclude all households that reside in a county with a radio station that broadcasts the *Dave Ramsey Show* at some point. Columns 5–6 apply both restrictions. Robust standard errors clustered at the zip code level and shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.7. Robustness: Log items – Excluding counties with affiliates and areas close to Nashville

	Dependent variable: log(Number of purchased products)					
	Drop zip codes close to Nashville		Drop counties with affiliate stations		Apply both restrictions	
	(1)	(2)	(3)	(4)	(5)	(6)
Radio coverage	-0.016***	-0.021***	-0.019***	-0.020***	-0.017***	-0.019***
	(0.003)	(0.004)	(0.004)	(0.005)	(0.004)	(0.005)
N	3,337,267	3,040,998	2,309,039	2,031,659	2,045,509	1,799,901
R^2	0.542	0.549	0.542	0.551	0.542	0.550
Mean of dep. var.	4.182	4.179	4.204	4.201	4.198	4.194
Household & Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
Household controls		Yes		Yes		Yes
Local economic conditions		Yes		Yes		Yes
State x Time FEs		Yes		Yes		Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in all regressions is the log of the number of purchased products per month. Columns 1–2 exclude households residing in zip codes within 500 km of Nashville, Tennessee. Columns 3–4 exclude all households that reside in a county with a radio station that broadcasts the *Dave Ramsey Show* at some point. Columns 5–6 apply both restrictions. Robust standard errors clustered at the zip code level and shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.8. Robustness: Availability of the radio show on other channels

	Excluding years after joining:		
	(1) 2016 SiriusXM	(2) 2015 Everydollar	(3) 2013 YouTube
Panel A: Log expenditures			
Radio coverage	-0.011***	-0.010***	-0.010***
	(0.003)	(0.003)	(0.003)
N	3,248,939	2,935,565	2,604,519
R^2	0.528	0.534	0.541
Mean of dep. var.	6.182	6.180	6.175
Household FEs	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
Panel B: Log items			
Radio coverage	-0.010***	-0.007**	-0.006**
	(0.003)	(0.003)	(0.003)
N	3,240,312	2,927,445	2,597,187
R^2	0.557	0.566	0.575
Mean of dep. var.	4.188	4.188	4.190
Household FEs	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in Panel A is the log of monthly household expenditures. The dependent variable in Panel B is the log of the number of purchased items. Columns exclude all observations after the point in time when the *Dave Ramsey Show* launched on the channel indicated in the column header. The radio show launch on SiriusXM in November 2016. It launched EveryDollar.com in March 2013. Robust standard errors clustered at the zip code level and shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.9. Robustness – Alternative clustering of standard errors

	Robust standard errors clustered at the level of:			
	(1) Zip code	(2) County	(3) DMA	(4) State
Panel A: Log expenditures				
Radio coverage	-0.0131***	-0.0131***	-0.0131***	-0.0131***
	(0.0027)	(0.0032)	(0.0036)	(0.0028)
N	3,744,066	3,744,066	3,744,066	3,744,066
R^2	0.518	0.518	0.518	0.518
Mean of dep. var.	6.185	6.185	6.185	6.185
Household FEs	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes
Panel B: Log items				
Radio coverage	-0.0168***	-0.0168***	-0.0168***	-0.0168***
	(0.0029)	(0.0035)	(0.0041)	(0.0038)
N	3,734,881	3,734,881	3,734,881	3,734,881
R^2	0.541	0.541	0.541	0.541
Mean of dep. var.	4.189	4.189	4.189	4.189
Household FEs	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates where the unit of observation is a household-month. The dependent variable in Panel A are log expenditures. The dependent variable in Panel B are log purchased items. Each column uses robust standard errors clustered at the geographic or administrative unit indicated by the column header.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.10. Robustness – Using Nielsen projection factors to re-weigh households

	(1)	(2)	(3)	(4)
Panel A: Log expenditures				
Radio coverage	-0.0110***	-0.0099**	-0.0156***	-0.0155***
	(0.0040)	(0.0039)	(0.0041)	(0.0051)
N	3,683,294	3,683,294	3,353,738	3,353,738
R^2	0.530	0.533	0.535	0.538
Mean of dep. var.	6.145	6.145	6.148	6.148
Household & Time FEs	Yes	Yes	Yes	Yes
Household controls		Yes	Yes	Yes
Local economic conditions			Yes	Yes
State x Time FEs				Yes
Panel B: Log items				
Radio coverage	-0.0169***	-0.0152***	-0.0232***	-0.0250***
	(0.0043)	(0.0042)	(0.0045)	(0.0054)
N	3,674,329	3,674,329	3,345,823	3,345,823
R^2	0.555	0.558	0.559	0.562
Mean of dep. var.	4.158	4.158	4.156	4.156
Household & Time FEs	Yes	Yes	Yes	Yes
Household controls		Yes	Yes	Yes

Notes: This table shows WLS regression estimates of equation 1.4.1. Households are weighted using the weights supplied by Nielsen. Households with weights above 10,000 are excluded. The dependent variable in Panel A are log expenditures. The dependent variable in Panel B are log purchased items. Robust standard errors clustered at the zip code level and shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

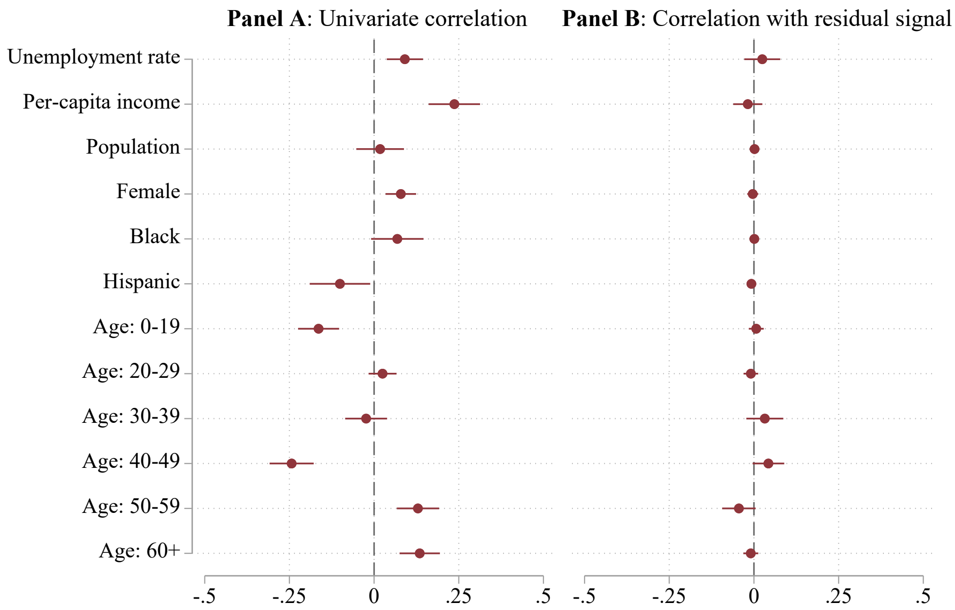

Figure 1.C.8. Residual signal strength and time-varying characteristics

Notes: This figure plots OLS regression coefficients on a county-year panel using different time-varying county-level characteristics as dependent variable. Dependent variables are standardized to have mean zero and standard deviation one to facilitate comparisons. Each point estimate is obtained from a separate regression. Panel A reports the regression coefficient between the time-varying county characteristics and the standardized, predicted radio signal strength. Panel B reports analogous estimates conditional on the predicted free-space signal, its square, county and year fixed effects, and region×year fixed effects. The county-year panel is derived from the baseline sample by collapsing variables to the county-year level. Robust standard errors clustered at the DMA level are used to construct 95% confidence intervals.

Table 1.C.11. Robustness: Heterogeneity analysis by financial struggles without controls

	Dependent variable: log (Expenditures)					
	Expenditures		Income		Expenditures to income	
	(1) High	(2) Low	(3) High	(4) Low	(5) High	(6) Low
Radio show	-0.020***	-0.005	-0.013***	-0.014***	-0.020***	-0.007*
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
N	1,982,051	1,762,015	2,032,501	1,711,565	1,864,927	1,879,139
R^2	0.453	0.445	0.518	0.513	0.518	0.476
Mean of dep. var.	6.449	5.889	6.233	6.129	6.353	6.019

Notes: This table uses 2004–2019 Nielsen Homescan data at the household-by-month level. The dependent variable is the log of household expenditures. “Radio show” is a binary indicator taking value one after local market entry of the *Dave Ramsey Show*. Robust standard errors clustered by zip code are shown in parentheses. Each column provides estimates from a subset of households obtained by a median split based on the household covariate indicated in the column’s header. For the median split in columns 1–2, I use the average, inflation-adjusted and equivalized expenditures in the first year in the panel. For the median split in columns 3–4, I use the average, inflation-adjusted and equivalized household income in the first year in the panel. For the median split in columns 5–6, I use the average household expenditures normalized by income in the first year in the panel.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 1.D Additional analyses

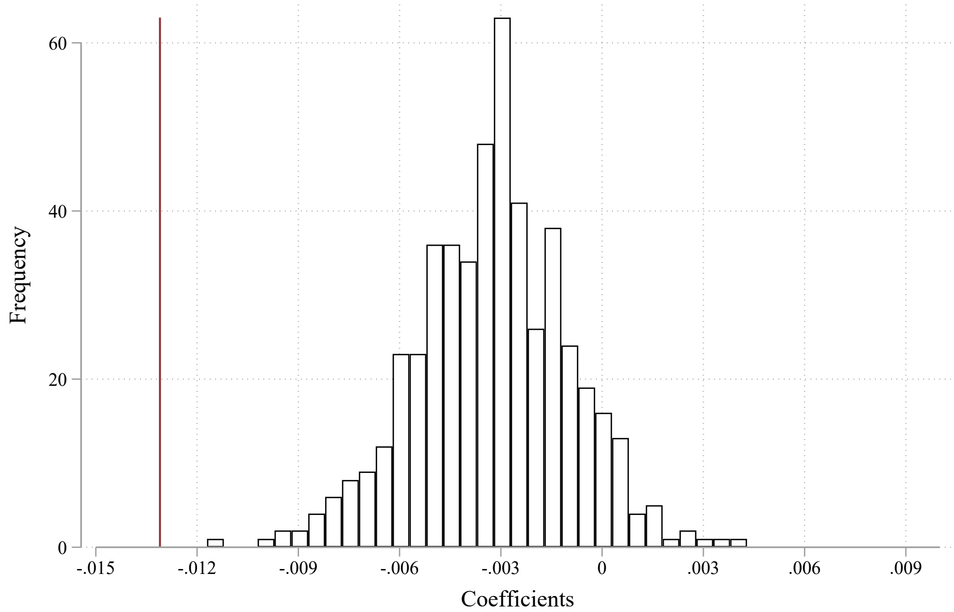

Figure 1.D.1. Counterfactual radio coverage

Notes: This figure plots a histogram of coefficients from a regression of log household expenditures on a counterfactual radio coverage indicator, including household and year-month fixed effects (see equation (1.4.1)). The distribution of coefficients is obtained from 500 counterfactual assignments. The coefficient from the actual radio coverage is shown as a red vertical line. Each counterfactual estimate is obtained as follows. I repeatedly assign a randomly chosen counterfactual market entry date to each zip code. If a zip codes is outside the actual coverage area of all affiliated radio stations, the zip code is always assigned to the control group without any market entry. I thus only vary the timing but not the set of zip codes that eventually receive access to the radio show. Based on the counterfactual timing of market entry, I apply equivalent sample restrictions as described in Section 1.3.2, and re-estimate equation 1.4.1 without time-varying controls using household expenditures as the dependent variable.

Table 1.D.1. Presidential elections: Turnout and voting behavior

	Turnout in Presidential election			Republican vote share		
	(1)	(2)	(3)	(4)	(5)	(6)
Radio coverage	-0.009	-0.005	0.003	0.008**	0.002	0.004
	(0.007)	(0.005)	(0.003)	(0.004)	(0.003)	(0.002)
N	15415	15415	15410	15415	15415	15410
R^2	0.937	0.963	0.977	0.943	0.954	0.977
Mean of dep. var.	0.470	0.470	0.471	0.563	0.563	0.563
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Baseline covar. x Year FEs		Yes	Yes		Yes	Yes
State x Year FEs			Yes			Yes

Notes: This table shows OLS regression estimates using electoral outcomes from the Presidential elections in 2000–2016. The unit of observation is a county-election. Turnout is measured as the ratio of cast votes to the voting age population. Radio coverage is the share of the county population with radio coverage. Observations are weighted by the voting age population Baseline county characteristics in 2000 include the percent of females, blacks, Hispanics and age group shares in 10 year bins. Robust standard errors clustered at the state level are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 1.E Content analysis

1.E.1 Qualitative evidence

This section contains a collection of quotes from Dave Ramsey that shed light on his views on consumption, debt and the role of social and cultural expectations.

Social and cultural expectations

- "We buy things we don't need with money we don't have to impress people we don't like."
- "We lived our lives according to the standards set to 'keep up with the Joneses.' Turns out they were broke and living in debt, too." (The Total Money Makeover, p. 20)
- "It is human nature to want it and want it now; it is also a sign of immaturity. Being willing to delay pleasure for a greater result is a sign of maturity. However, our culture teaches us to live for the now. 'I want it' we scream, and we can get it if we are willing to go into debt. Debt is a means to obtain the 'I want its' before we can afford them." (The Total Money Makeover, p. 16)
- "We live in a culture that quit asking, 'How much?' and instead asks, 'How much down, and how much a month?'" (The Total Money Makeover, p. 33)
- "Peer pressure, cultural expectations, 'reasonable standard of living' – I don't care how you say it, we all need to be accepted by our crowd and our families. This need for approval and respect drives us to do some really insane things. One of the paradoxically dumb things we do is to destroy our finances by buying garbage we can't afford to try to make ourselves appear wealthy to others." (The Total Money Makeover, p. 78)
- "Peer pressure is very, very powerful. 'We are scaling down' is a painful statement to make to friends or family. 'We will have to pass on that trip or dinner because it is not in our budget' is virtually impossible for some people to say. Being real takes tremendous courage. We like approval, and we like respect, and to say otherwise is another form of denial. The wish for the admiration of others is normal. The problem is that this admiration can become a drug. Many of you are addicted to this drug, and the destruction to your wealth and financial well-being caused by your addiction is huge." (The Total Money Makeover, p. 80)
- "Financial peace isn't the acquisition of stuff. It's learning to live on less than you make, so you can give money back and have money to invest. You can't win until you do this."
- "You must walk to the beat of a different drummer. The same beat that the wealthy hear. If the beat sounds normal, evacuate the dance floor immediately! The goal is to not be normal, because as my radio listeners know, normal is broke."

- "70% of Americans live paycheck to paycheck. Seven out of ten people you walk past going down the sidewalk are broke. You can model your life after them, and you will be one of them. Or you can mode your life after the weird people. Because wealth is unusual. It's not normal. So you have to engage in unusual behaviors and habits to create unusual results."

Debt

- "Debt has been sold to us so aggressively, so loudly, and so often that to imagine living without debt requires myth-busting." (The Total Money Makeover, p. 17)
- "Debt is so ingrained into our culture that most Americans cannot even envision a car without a payment, a house without a mortgage, a student without a loan, and credit without a card." (The Total Money Makeover, p. 17-18)
- "Debt is not a tool; it is a method to make banks wealthy, not you. The borrower truly is slave to the lender." (The Total Money Makeover, p. 48)
- "My contention is that debt brings on enough risk to offset any advantages that could be gained through leverage of debt." (The Total Money Makeover, p. 20)
- "Larry Burkett said debt is not the problem; it is the symptom. I feel debt is the symptom of overspending and undersaving." (The Total Money Makeover, p. 45)

Behavior

- "Winning at money is 80 percent behavior and 20 percent head knowledge. What to do isn't the problem; doing it is. Most of us know what to do, but we just don't do it. If I can control the guy in the mirror, I can be skinny and rich." (The Total Money Makeover, p. 3)
- "I teach concepts, not mathematical formulas." (The Total Money Makeover, p. xvi)
- "Break through the temptation to remain in the same situation, and opt for the pain of change before the pain of not changing searches you out." (The Total Money Makeover, p. 14)
- "Living on less than you make is matter of controlling yourself, not a matter of math."
- "I can always tell which ones are serious and which aren't. There's something in their voices that communicates passion and conviction when they're really excited about getting out of debt. But if they're just playing around with the idea, if they're simply curious about it, then their voices are flat. If I don't hear any passion behind what they're saying, I know they aren't ready to cut up the credit cards and dump their debt for good. That's because getting out of debt isn't about solving a math problem; it's about changing your life–and that requires a change of heart."

- "One thing I am sure of in my Total Money Makeover: I had to quit telling myself that I had innate discipline and fabulous natural self-control. That is a lie. I have to put systems and programs in place that make me do smart things. Saying, 'Cross my fingers and hope to die, I promise, promise, promise I will pay extra on my mortgage because I am the one human on the planet who has that kind of discipline,' is kidding yourself. A big part of being strong financially is that you know where you are weak and take action to make sure you don't fall prey to the weakness."

1.E.2 Quantitative evidence

1.E.2.1 Data and text processing

I use a Python-based command line program to collect the automatically generated subtitles of the 5,587 YouTube videos uploaded by the *Dave Ramsey Show* between August 13, 2013, and May 31, 2021. These subtitles are available in `WebVTT` format, which includes both the audio transcripts as well as timestamps indicating the start time for each line of text. I remove timestamps and aggregate subtitles to documents containing 5 minutes of contiguous speech.

I apply a series of commonly used processing steps to prepare the raw text data for analysis. I convert the text to lowercase and remove whitespace. Next, I remove English language stopwords that occur very frequently. In addition, I remove numerals (e.g. "five", "thousand") as those occur frequently when Dave Ramsey asks callers for information about their finances. Moreover, I remove a list context specific words mentioned in the radio show's jingle and during commercial breaks: `headquarter`, `bmw`, `king`, `blinds.com`, `promo`, `code`, `sample`, `churchill`, `zander`, `mama`, `ship`, `shipping`, `blinds`, `window`, `special`, `smartvestor`. I also remove names of personalities appearing on the radio show such as `dave`, `ramsey`, `chris` and `logan`. I then apply the Porter stemmer, one of the most common English language stemming algorithms. I remove all non-alphanumeric characters, exclude words that occur less than 100 times, and all words that only include numbers.

1.E.2.2 Latent Dirichlet Allocation

For topic analysis, I use Latent Dirichlet Allocation (LDA) which is an unsupervised machine learning technique for topic modeling (Gentzkow, Kelly, and Taddy, 2019). As an input, I use the document-term matrix of all unigrams that appear in at most 90% of all documents. Here, a document corresponds to the words spoken in a 5-minute interval. I train the LDA model using an online learning method with hyperparameters $\kappa = 0.7$, $\tau_0 = 10$ and a batch size of 512.

Figure 1.E.1 shows the 50 words with the highest probability by topic. Differences in the size corresponds to differences in probabilities. To assign labels to topics, I rely both on the word cloud and manual inspection of text segments where the model has a high confidence in its classification.

1.E.2.3 Keywords and word co-occurrences

To complement the topic model approach, I explore common words and their associations across documents. Table 1.E.1 provides an overview of the most frequent words and keywords across the 5,587 YouTube videos uploaded by the *Dave Ramsey Show*. Figure 1.E.2 illustrates the network of words with the highest co-occurrence rates with the word "debt", using a methodology proposed by Bail (2016) and ex-

Figure 1.E.1. Word distribution by topic

Notes: This figure illustrates the results of the Latent Dirichlet Allocation. Each word cloud shows the top 50 words by topic based on their frequency of occurrence. The font size is proportional to the word frequency.

cluding the same set of stop words as in the LDA analysis. This complementary approach confirms that paying off debt is a central theme of the radio show.

Table 1.E.1. Most frequent spoken words and keywords for YouTube videos of the *Dave Ramsey Show*

Rank	Word	Frequency	Video keyword	Frequency
1	money	105286	money	2553
2	debt	78079	credit card	2546
3	pay	64709	real estate	2544
4	start	60781	buy	2518
5	dollar	56420	insurance	2510
6	hous	51554	save	2509
7	car	47284	how to make money	2504
8	life	47121	snowball	2501
9	live	46014	buying house	2497
10	month	45808	compound interest	2495
11	save	38015	budget money debt cash	2493
12	busi	34036	debt	2213
13	home	33057	debt free scream	868
14	loan	31505	personal finance	697
15	incom	31105	budget	559
16	paid	30377	student loans	369
17	job	30140	finance	355
18	plan	29521	drtlgi	336
19	step	29282	family	331
20	question	28915	credit	326
21	buy	28783	marriage	316
22	free	28470	investing	301
23	love	27938	debt free	293
24	stuff	26490	paying off debt	266
25	kid	26236	free	247
26	financi	25822	loans	244
27	fund	25414	student loan debt	236
28	famili	25327	loan	235
29	care	23960	car	227
30	babi	23864	scream	214
31	budget	23412	pay off debt	213

Notes: This table shows 30 most frequent spoken words as well as the most commonly used keywords attached to YouTube videos uploaded by the channels “The Ramsey Show – Full Episodes” and “The Ramsey Show – Highlights” between August 13, 2013, and May 31, 2021. The list excludes all keywords that include “dave”, “ramsey”, “video”, or “show”. The most frequent spoken words exclude a list of commonly used English words.

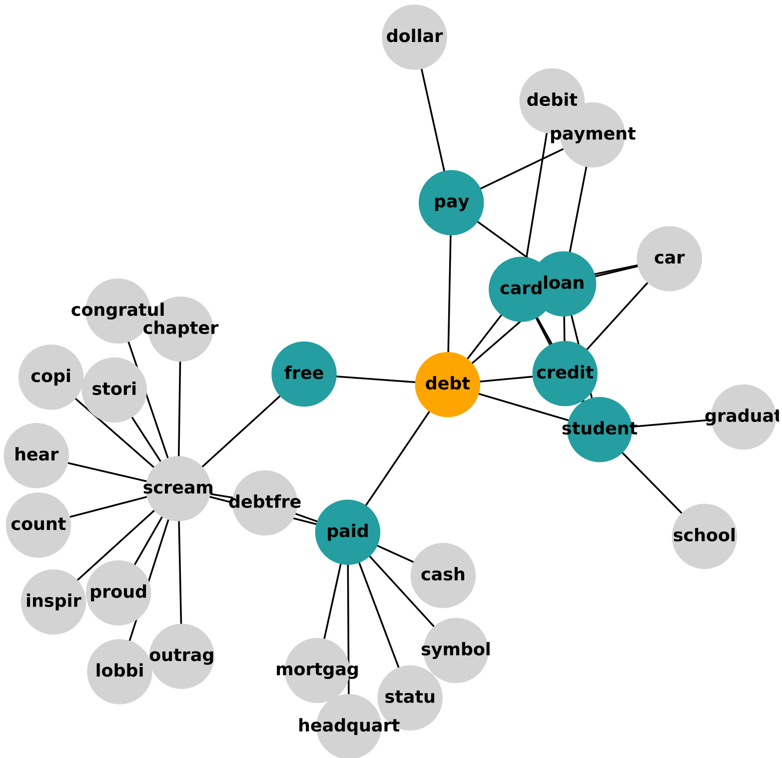

Figure 1.E.2. Correlates of the word "debt"

Notes: This figure uses data from YouTube. It shows words that frequently occur with the word "debt" in a 5 minute segment of audio. Edges between words indicate that when constructing binary indicators for the presence of these words in a document, these indicators have a correlation of 0.20 or above. To generate this list, I start with the word "debt" and collect all words with a correlation of at least 0.20. For these "direct links", I obtain all words that have a correlation of at least 0.30. I then plot the connections among these words.

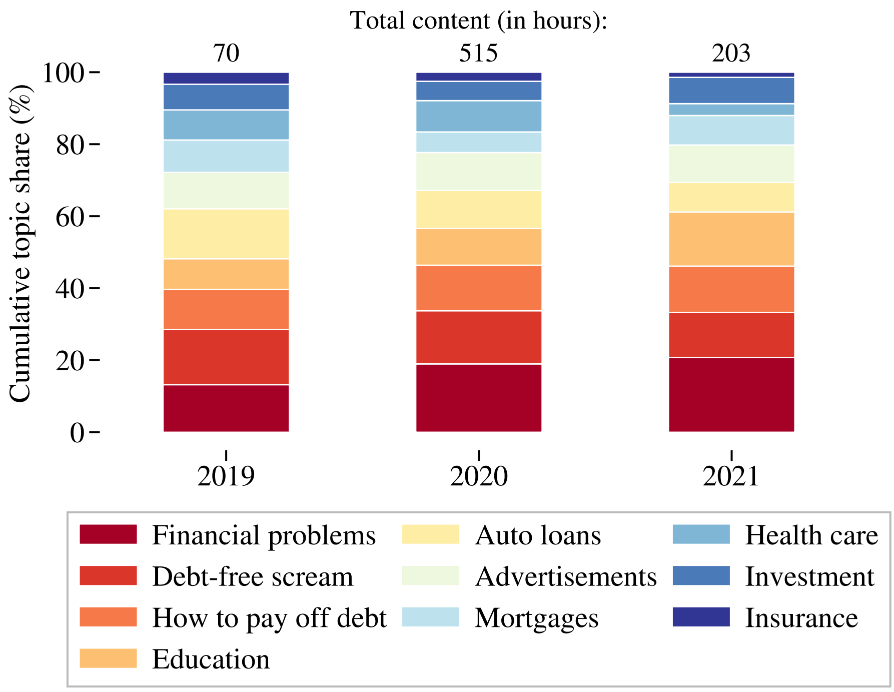

Figure 1.E.3. Topic distribution: Restricting to full episodes

Notes: This figure displays the distribution of topics featured in the *Dave Ramsey Show* in the videos uploaded on its YouTube channel. This figure restricts to videos covering full episodes of the show. Topic shares are obtained from *Latent Dirichlet Allocation* by calculating the average probability of each topic across documents. For each year, the total content (in hours) uploaded on YouTube is indicated above each bar.

Appendix 1.F Experiment

This section contains additional material and information about the survey experiment discussed in Section 1.6.

1.F.1 Research transparency

Preregistration The main experiment was preregistered on the AEA RCT Registry as project #AEARCTR-0008050. The preregistration includes details on the experimental design, the sampling process, planned sample size, exclusion criteria, hypotheses and the main analyses. Below, I document deviations from the preregistration:

- The preregistration uses a different title and different treatment labels.
- The preregistration did not include quotas based on sociodemographic characteristics. In practice, the sampling process was stratified based on age, gender and education, which results in a more representative sample of the US population.
- Respondents below the age of 18 and those who do not reside in the US were not eligible to participate in the survey, which was not preregistered.
- When construction attitudinal indices, I normalize the indices using the mean and standard deviation in the control group used in the analysis. The preregistration did not specify the reference group for the normalization. However, the normalization does not affect the economical or statistical significance of the results.
- In contrast to the preregistration, I include control variables when estimating treatment effects in the main experiment. The results are robust to not including controls, as shown in Table 1.F.5.
- Non-preregistered analyses include (i) a robustness exercises estimating treatment effects on individual items used to measure attitudes and (ii) the descriptive evidence on the correlation between attitudes and behavior.

The one-week follow-up survey was not preregistered.

Ethical approval The experimental study received ethics approval from the German Association for Experimental Economic Research (#T7wapLjB, 07/20/2021).

Data and code availability The experimental data and the analysis code will be made available online.

Competing interests I declare no competing interests.

1.F.2 Figures and Tables

Panel A: Main experiment – Design overview

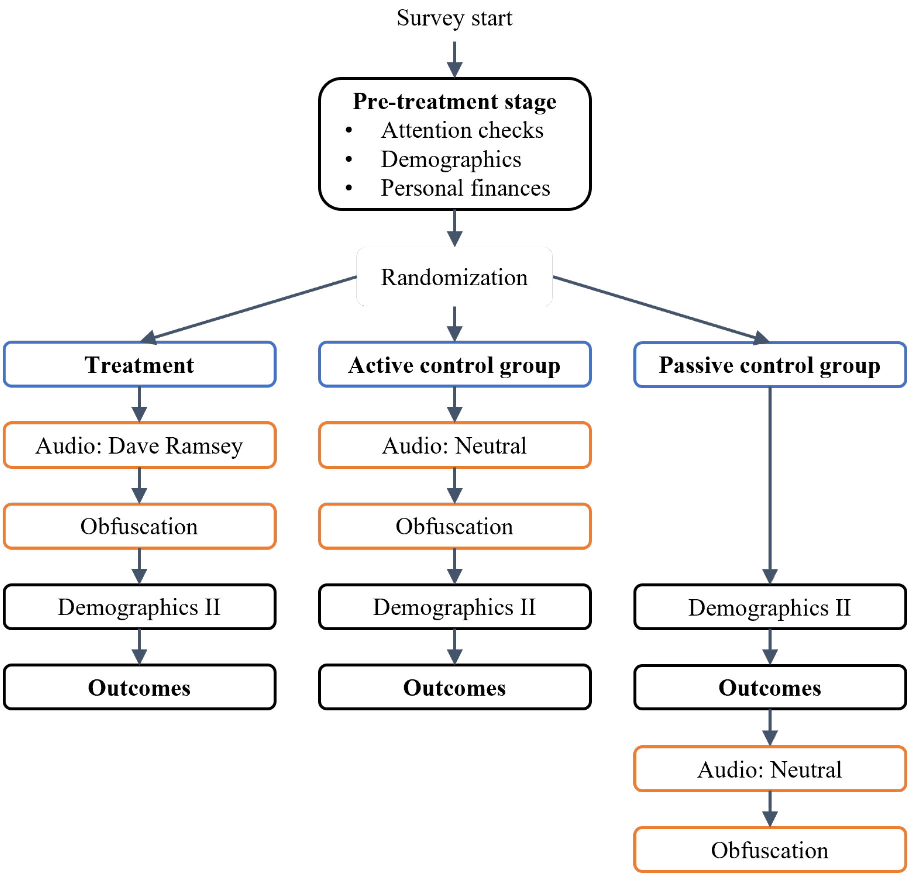

Panel B: Timing of the main experiment and the follow-up survey

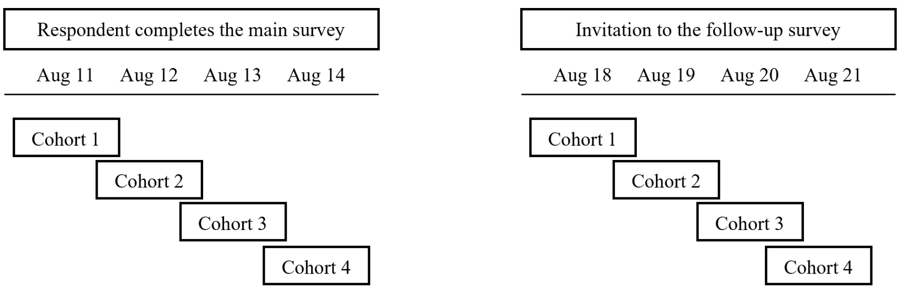

Figure 1.F.1. Design overview and timing

Notes: This figure provides and overview of the main design features used in the mechanism experiment.

Table 1.F.1. Comparison of the survey sample to the general US population

Variable	Survey sample	American Community Survey (2019)
Female	50%	51%
Age: 18–34	30%	30%
Age: 35–54	30%	32%
Age: 55+	40%	38%
Education: Bachelor’s degree or above	30%	31%
Region: Northeast	19%	17%
Region: Midwest	21%	21%
Region: South	43%	38%
Region: West	17%	24%

Notes: This table provides summary statistics for the sample in the main experiment (column 1) and the general US population (column 2) for basic demographic characteristics.

Table 1.F.2. Test of balance: Main experiment

	Means (std. dev.)			Differences (*p*-values)		
	Treatment group (T)	Active control (A)	Passive control (P)	T - A	T - P	A - P
Age	47.825	48.015	48.071	-0.190	-0.245	0.056
	(17.763)	(17.504)	(18.351)	(0.868)	(0.828)	(0.961)
Female	0.494	0.504	0.491	-0.010	0.003	-0.014
	(0.500)	(0.501)	(0.500)	(0.749)	(0.918)	(0.668)
College degree	0.445	0.447	0.446	-0.002	-0.001	-0.001
	(0.497)	(0.498)	(0.498)	(0.958)	(0.975)	(0.982)
Log income	10.628	10.558	10.646	0.070	-0.018	0.088
	(0.891)	(0.930)	(0.889)	(0.232)	(0.750)	(0.126)
Log debt	6.302	6.170	6.186	0.133	0.117	0.016
	(4.538)	(4.461)	(4.539)	(0.647)	(0.680)	(0.954)
Democrat	0.437	0.417	0.429	0.020	0.008	0.012
	(0.497)	(0.494)	(0.495)	(0.532)	(0.805)	(0.693)
Republican	0.297	0.285	0.283	0.012	0.014	-0.003
	(0.457)	(0.452)	(0.451)	(0.691)	(0.616)	(0.928)
Subjective financial literacy	4.699	4.523	4.619	0.176*	0.080	0.096
	(1.405)	(1.452)	(1.295)	(0.057)	(0.341)	(0.270)
Savings ability	0.638	0.587	0.608	0.051	0.030	0.021
	(0.481)	(0.493)	(0.489)	(0.105)	(0.315)	(0.507)
Region: Northeast	0.222	0.160	0.178	0.062**	0.043*	0.019
	(0.416)	(0.367)	(0.383)	(0.015)	(0.084)	(0.426)
Region: Midwest	0.205	0.191	0.242	0.014	-0.036	0.050*
	(0.404)	(0.394)	(0.428)	(0.592)	(0.163)	(0.055)
Region: South	0.396	0.472	0.416	-0.076**	-0.020	-0.056*
	(0.490)	(0.500)	(0.493)	(0.017)	(0.514)	(0.074)
Region: West	0.177	0.177	0.164	0.000	0.013	-0.013
	(0.382)	(0.382)	(0.370)	(0.992)	(0.572)	(0.583)
p-value of joint *F*-test				0.313	0.796	0.689
Observations	492	470	538	962	1,030	1,008

Notes: This table shows a test of balance for the main experiment. Columns 1–3 show the means and standard deviations of respondent covariates in the different treatments arms. Columns 4–6 show differences in means between the groups indicated in the column header together with *p*-values in parentheses. The *p*-values of the joint *F*-test are determined by regressing the treatment indicator on the vector of covariates. The *F*-test tests the joint hypothesis that none of the covariates predicts treatment assignment.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.3. Balance of post-treatment demographics

	Means (std. dev.)			Differences (p-values)		
	Treatment group (T)	Active control (A)	Passive control (P)	T - A	T - P	A - P
Black	0.126	0.126	0.113	0.000	0.013	-0.012
	(0.332)	(0.332)	(0.317)	(0.982)	(0.533)	(0.553)
White	0.799	0.777	0.805	0.022	-0.006	0.028
	(0.401)	(0.417)	(0.397)	(0.401)	(0.808)	(0.271)
Hispanic	0.083	0.074	0.072	0.009	0.011	-0.002
	(0.277)	(0.263)	(0.260)	(0.611)	(0.517)	(0.905)
Full-time employment	0.325	0.338	0.325	-0.013	-0.000	-0.013
	(0.469)	(0.474)	(0.469)	(0.667)	(0.998)	(0.662)
Unemployed	0.108	0.111	0.126	-0.003	-0.019	0.016
	(0.310)	(0.314)	(0.333)	(0.885)	(0.353)	(0.441)
Not in labor force	0.376	0.360	0.348	0.016	0.028	-0.012
	(0.485)	(0.480)	(0.477)	(0.598)	(0.343)	(0.691)
p-value of joint F-test				0.837	0.816	0.865
Observations	492	470	538	962	1,030	1,008

Notes: This table shows a balance test for the main experiment using post-treatment demographic variables. Columns 1–3 show the means and standard deviations of respondent covariates in the different treatments arms. Columns 4–6 show differences in means between the groups indicated in the column header together with p-values in parentheses. The p-values of the joint F-test are determined by regressing the treatment indicator on the vector of covariates. The F-test tests the joint hypothesis that none of the covariates predicts treatment assignment.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.4. Correlation between attitudes and past behavior

	Log debt		Debt-free		Log spending	
	(1)	(2)	(3)	(4)	(5)	(6)
Debt attitudes	0.391**	0.398***	-0.045***	-0.047***	-0.049	-0.037
	(0.152)	(0.151)	(0.015)	(0.015)	(0.049)	(0.046)
Consumption attitude	0.171	0.213	-0.015	-0.015	0.169***	0.107***
	(0.145)	(0.152)	(0.015)	(0.016)	(0.042)	(0.039)
N	1,008	1,008	1,008	1,008	1,008	1,008
Mean of dep. var.	6.178	6.178	0.301	0.301	4.805	4.805
Controls	No	Yes	No	Yes	No	Yes

Notes: This table shows OLS regression estimates using respondents from the main study, excluding respondents in the treatment group. The debt attitude index and the consumption attitude index are constructed as described in the main text and oriented such that larger values correspond to more positive attitudes towards the object. Both indices are normalized to have mean zero and standard deviation one. Control variables include numerical age and age squared, log income, female indicator, and an indicator for having completed a Bachelor's degree or higher.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.5. Robustness: Treatment effects on attitudes across studies without controls

	Main study		Robustness: Passive control		One-week follow-up	
	(1) Debt attitudes	(2) Consumption attitudes	(3) Debt attitudes	(4) Consumption attitudes	(5) Debt attitudes	(6) Consumption attitudes
Treatment	-0.535***	-0.219***	-0.605***	-0.227***	-0.313***	-0.187*
	(0.065)	(0.065)	(0.061)	(0.065)	(0.096)	(0.099)
N	962	962	1,030	1,030	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	No	No

Notes: This table shows OLS regression estimates where the dependent variables are attitudes towards consumption and debt. The debt attitude index and the consumption attitude index are constructed as described in the main text and oriented such that larger values correspond to more positive attitudes towards the object. Both indices are normalized to have mean zero and standard deviation one. "Treatment" is a binary indicator taking value one for respondents who listened to a five-minute recording from the *Dave Ramsey Show*. Columns 1 and 2 use respondents from the main study assigned to the treatment group or the control group. Columns 3 and 4 use respondents from the main study assigned to the treatment group or the robustness control group. Columns 5 and 6 use respondents from the one-week follow-up survey pooling respondents from both control group conditions as a joint control group. Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.6. Main experiment – Treatment effects on attitudes by item

	Debt attitudes				Consumption attitudes	
	(1) There is no excuse for borrowing money	(2) You should always save up first before buying something	(3) You can live a good life without borrowing money	(4) All in all, borrowing money is not worth the cost	(5) I admire people who own expensive homes, cars, and clothes	(6) The things I own say a lot about how well I'm doing in life
Panel A: Active control						
Treatment	0.318***	0.270***	0.363***	0.507***	-0.134**	-0.257***
	(0.066)	(0.062)	(0.061)	(0.064)	(0.064)	(0.066)
N	962	962	962	962	962	962
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Passive control						
Treatment	0.452***	0.292***	0.352***	0.590***	-0.221***	-0.176***
	(0.065)	(0.059)	(0.057)	(0.061)	(0.063)	(0.064)
N	1,030	1,030	1,030	1,030	1,030	1,030
z-scored	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates using respondents from the main experiment. The dependent variables are respondents' agreement with the statements indicated in the column header and measured on a 5-point Likert scale from "Strongly agree" to "Strongly disagree". Responses are coded such that larger values indicate stronger agreement, and z-scored using the mean and standard deviation in the respective control group. "Treatment" is a binary indicator taking value one for respondents who listened to a five-minute recording from the *Dave Ramsey Show*. Panel A uses respondents from the treatment group and the control group. Panel B uses respondents from the treatment group and the robustness control group. Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.7. Follow-up survey – Treatment effects on attitudes by item

	Debt attitudes				Consumption attitudes	
	(1) There is no excuse for borrowing money	(2) You should always save up first before buying something	(3) You can live a good life without borrowing money	(4) All in all, borrowing money is not worth the cost	(5) I admire people who own expensive homes, cars, and clothes	(6) The things I own say a lot about how well I'm doing in life
Panel A: Baseline						
Treatment	0.277***	0.035	0.339***	0.260***	-0.123	-0.211**
	(0.093)	(0.101)	(0.093)	(0.096)	(0.094)	(0.100)
N	522	522	522	522	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Controls						
Treatment	0.260***	0.030	0.323***	0.267***	-0.142*	-0.228**
	(0.091)	(0.100)	(0.093)	(0.096)	(0.086)	(0.095)
N	522	522	522	522	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Panel C: IPAW						
Treatment	0.275***	0.049	0.361***	0.266***	-0.124	-0.215**
	(0.093)	(0.103)	(0.094)	(0.097)	(0.094)	(0.100)
N	522	522	522	522	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes
Panel D: Controls & IPAW						
Treatment	0.259***	0.041	0.344***	0.274***	-0.145*	-0.232**
	(0.091)	(0.102)	(0.094)	(0.097)	(0.086)	(0.095)
N	522	522	522	522	522	522
z-scored	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows regression estimates using respondents from the one-week follow-up survey. The dependent variable are respondents' agreement with the statements indicated in the column header and measured on a 5-point Likert scale from "Strongly agree" to "Strongly disagree". Responses are coded such that larger values indicate stronger agreement, and z-scored using the mean and standard deviation of non-treated respondents. "Treatment" is a binary indicator taking value one for respondents who listened to a five-minute recording from the *Dave Ramsey Show*. Panel A presents baseline OLS estimates without controls . Panel B includes numerical age and age squared, log income, female indicator, an indicator for having completed a Bachelor's degree or higher, and region indicators as controls. Panel C uses inverse probability of attrition weights (IPAW) obtained from a logistic regression of the attrition status dummy on the vector of baseline covariates from Table 1.F.2 to reweigh respondents. Panel D adds the control variables from Panel B to the specification from Panel C. Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.8. Follow-up survey – Test for differential attrition across treatment arms

	(1) Age	(2) Female	(3) College	(4) Log income	(5) Log debt	(6) Democrat	(7) Republican	(8) Financial literacy	(9) Savings ability	(10) Northeast	(11) Midwest	(12) South	(13) West
Treatment	-0.40	0.01	-0.03	0.01	0.45	0.03	0.01	0.04	0.03	0.05*	0.00	-0.05	-0.00
	(1.21)	(0.03)	(0.03)	(0.06)	(0.30)	(0.03)	(0.03)	(0.10)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Follow-up	4.04***	0.01	-0.06*	-0.10*	0.63**	0.03	0.03	0.04	-0.05	-0.01	0.01	0.03	-0.04
	(1.18)	(0.03)	(0.03)	(0.06)	(0.30)	(0.03)	(0.03)	(0.09)	(0.03)	(0.02)	(0.03)	(0.03)	(0.02)
Treatment x Follow-up	0.54	-0.02	0.09	0.05	-0.94*	-0.04	-0.00	0.24	0.03	-0.00	-0.05	0.02	0.03
	(2.04)	(0.06)	(0.06)	(0.10)	(0.53)	(0.06)	(0.05)	(0.16)	(0.06)	(0.05)	(0.05)	(0.06)	(0.04)
Constant	46.64***	0.49***	0.47***	10.64***	5.96***	0.41***	0.27***	4.56***	0.62***	0.17***	0.21***	0.43***	0.18***
	(0.70)	(0.02)	(0.02)	(0.04)	(0.18)	(0.02)	(0.02)	(0.05)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
N	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500

Notes: This table shows OLS regression estimates using baseline demographic characteristics as dependent variable. Each regression includes the full interaction between the binary treatment indicator and a binary dummy indicating whether a respondent is part of the follow-up sample. Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.F.9. Test of balance: Follow-up survey

	Means (std. dev.)		Difference (p-values)
	Treatment group (T)	Control group (C)	T - C
Age	50.813	50.678	0.135
	(17.455)	(17.723)	(0.935)
Female	0.485	0.504	-0.019
	(0.501)	(0.501)	(0.686)
College degree	0.462	0.407	0.055
	(0.500)	(0.492)	(0.237)
Log income	10.595	10.538	0.057
	(0.822)	(0.888)	(0.483)
Log debt	6.097	6.586	-0.489
	(4.703)	(4.436)	(0.247)
Democrat	0.427	0.442	-0.015
	(0.496)	(0.497)	(0.751)
Republican	0.316	0.305	0.011
	(0.466)	(0.461)	(0.800)
Subjective financial literacy	4.877	4.598	0.279**
	(1.261)	(1.355)	(0.024)
Savings ability	0.626	0.564	0.062
	(0.485)	(0.497)	(0.181)
Region: Northeast	0.216	0.165	0.051
	(0.413)	(0.372)	(0.156)
Region: Midwest	0.181	0.225	-0.044
	(0.386)	(0.418)	(0.250)
Region: South	0.433	0.464	-0.032
	(0.497)	(0.499)	(0.497)
Region: West	0.170	0.145	0.024
	(0.376)	(0.353)	(0.471)
p-value of joint F-test			0.246
Observations	171	351	522

Notes: This table shows a test of balance for the sample in the follow-up survey. Columns 1–2 show the means and standard deviations of respondent covariates in the treatment group and the pooled control group comprising respondents in the control group and the robustness control group. Columns 3 show differences in means between the treatment group and the control group together with p-values in parentheses. The p-value of the joint F-test is determined by regressing the treatment indicator on the vector of covariates. The F-test tests the joint hypothesis that none of the covariates predicts treatment assignment.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.F.3 Secondary outcomes

To obfuscate the purpose of the main study, the survey includes an obfuscation module with several non-attitudinal measures. While these measures are not the primary interest of the experiment, this section provides a discussion of the treatment effects on these secondary outcomes. Table 1.F.10 presents estimates of the treatment effect of listening to the *Dave Ramsey Show* for five minutes on these outcomes using the audio control group (Panel A) or the robustness control group (Panel B) as comparison group.

First, column 1 shows that there is no statistically significant and robust treatment effect on respondents' demand for information about personal finances. Consistent with the hypothesis in the preregistration, the point estimate for the effect is larger when using the audio control group as comparison group, although the difference is not statistically significant at conventional levels. Second, consistent with my preregistered hypothesis, there is no statistically significant and robust treatment effect on general financial literacy as measured by the Big 5 survey module (column 2). Indeed, the audio recording from the *Dave Ramsey Show* does not include any information that would be help respondents answer the factual questions in the Big 5 module. Third, while I do not find treatment effects on respondents' beliefs about the average debt of US households (column 3), column 4 shows that treated respondents think that a larger share of Americans has any kind of debt ($p < 0.05$). The effect size is modest and depends on the comparison group and varies from 2.8 to 5.5 percentage points relative to a baseline of about 60-63%.

This provides further suggestive evidence that the *Dave Ramsey Show* affects the behavior of its listeners primarily by changing attitudes towards consumption and debt using its consistent and persuasive narrative.

Table 1.F.10. Treatment effects on secondary outcomes

	(1) Information demand	(2) Financial literacy	(3) Belief: Average debt	(4) Belief: Any debt
Panel A: Audio control group				
Treatment	0.052*	0.159*	3.231	5.457***
	(0.029)	(0.082)	(3.463)	(1.404)
Constant	0.253***	3.034***	75.376***	60.223***
	(0.020)	(0.058)	(2.337)	(1.067)
N	962	962	962	962
Panel B: Robustness control group				
Treatment	-0.004	-0.086	6.047*	2.841**
	(0.029)	(0.080)	(3.266)	(1.267)
Constant	0.309***	3.279***	72.560***	62.840***
	(0.020)	(0.055)	(2.034)	(0.880)
N	1,030	1,030	1,030	1,030

Notes: This table shows OLS regression estimates using respondents from the main experiment. "Information demand" takes value one for respondents who said that they would like to receive information about personal finances, and zero otherwise. "Financial literacy" is the number of correctly answered questions (out of 5) from the Big 5 financial literacy questionnaire. "Belief: Average debt" is the respondent's belief about the average debt of US Americans in thousand US dollars. "Belief: Any debt" is the belief about the share of Americans that have any debt at all. "Treatment" is a binary indicator taking value one for respondents who listened to the five-minute recording from the *Dave Ramsey Show*. Regressions do not include any control variables Panel A uses respondents from the treatment group and the control group (that listened to a neutral audio). Panel B uses respondents from the treatment group and the robustness control group (that did not listen to an audio recording). Robust standard errors are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.F.4 Experimental instructions: Main study

1.F.4.1 Pre-treatment stage

Welcome!

Thank you for your interest in completing this survey. The survey has **two parts** and takes about **15 minutes** to complete. By completing this survey, you help us understand how people in the US think about important questions. It is part of a study conducted by researchers from the University of Bonn.

You are not allowed to participate in this study more than once. If you experience a technical error or problem, do not try to restart or retake the study. Rather, send us an email with a description of your problem and we will get back to you. If you have any questions regarding this study, please email felix.chopra@uni-bonn.de

To participate in the study, you have to live in the US, and be 18 years or older.

[Page break]

Please consent to the processing of your data and our privacy policy *Click **here** to display the full privacy policy.*

Your data will be stored and analyzed in full compliance with the highest standards of the data protection laws of the European Union. In particular, no conclusions about your person will be drawn. You can withdraw your consent at any time.

- I consent
- I do not consent

[Page break]

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This compromises the results of research studies**. To show that you are reading the survey carefully, please choose both "Very strongly interested" and "Not at all interested" as your answer to the next question.**

Given the above, how interested are you in politics?

- Very strongly interested
- Very interested

- A little bit interested
- Not very interested
- Not at all interested

[Page break]

As part of this survey, you will listen to an audio recording. You can only participate in this survey if your device can play audio recordings. To see if this works, please try to play the audio below.

[Audio player with controls, see Figure 1.F.3]

Which color was mentioned in the audio recording?
[Dropdown menu]

[Page break]

Please provide us with some information about yourself.

What is your age?
[Dropdown menu]

What is your gender?

- Male
- Female
- Other / Prefer not to say

What was your annual gross household income in 2019?
[Dropdown menu]

What is the highest level of education you have completed or the highest degree you have received?

- Some high school, but no degree
- High school degree (or GED)
- Some college, but no degree
- Associate degree (2-year)
- Bachelor's degree (4-year)
- Post-graduate degree

With which political party do you identify the most?

- Democratic Party
- Republican Party
- Independent

[Page break]

How would you describe your overall financial knowledge?
[Very low (1), 2, 3, 4, 5, 6, Very high (7)]

Do you usually have money left over at the end of the month that you can save for larger purchases, emergency expenses or to build up savings?
[Yes, No]

Which, if any, of the following types of debt do you have? Please check all that apply.

- Mortgage debt
- Student loan debt
- Credit card debt
- Auto loan debt
- Other types of debt
- I have no debt

[if respondent did not select "I have no debt" in the previous question, display:]

In total, how much debt do you currently have?
[Dropdown menu]

What is the combined dollar value of all your spending on the categories below over the last 7 days?

- food consumed at home
- food consumed away from home
- leisure activities such as visiting the cinema or sport games
- clothing

The combined dollar of my spending on these categories over the last 7 days is...
[Text entry field]

[Page break]

We will now begin with the first part of this survey.

[Page break]

1.F.4.2 Treatments

On the next page, you will listen to a **5 minute** recording.

[Page break]

Please listen to this audio. We will ask you a few questions about it afterwards.
[Audio player with controls]

You will be able to advance to the next page once you finished listening to the audio.
[Page submit is visible after 5 minutes]

1.F.4.3 Obfuscation

Please answer these questions about the audio content you just listened to.

Did you enjoy listening to the content?
[Yes, No]

Imagine a local radio station near you would feature content like this. Would you be more or less likely to listen to this station?

- Much more likely
- Somewhat more likely
- About the same
- Somewhat less likely
- Much less likely

How would you rate the production quality of the content?

- Very high
- High
- Low

- Very low

How would you rate the novelty of the content?

- Very high
- High
- Low
- Very low

What is the name of the radio show that you just listened to?
[Text entry field]

On how many days do you listen to the radio in a typical week?
[Dropdown menu, values from 1 to 7]

Which, if any, of the following radio programs have you listened to in the past? Please select all that apply.

- Savage Nation
- Sean Hannity Show
- Dave Ramsey Show
- Marketplace
- BBC World Service
- Howard Stern Show
- Mark Levin Show
- Coast to Coast
- Morning Edition
- I don't listen to these radio shows

[Page break]

You will now continue to the second and final part of this survey.

[Page break]

Please answer these questions about yourself.

Which of the following best describes your race or ethnicity?
[Dropdown menu]

Are you of Hispanic, Latino, or Spanish origin?
[Yes, No]

What is your current employment status?
[Dropdown menu]

In which state do you currently reside?
[Dropdown menu]

What is your zipcode of residence?
[Text entry field]

1.F.4.4 Post-treatment measures

Suppose you had $100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow?

- More than $102
- Exactly $102
- Less than $102

Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account?

- More than today
- Exactly the same
- Less than today

If interest rates rise, what will typically happen to bond prices?

- They will rise
- They will fall
- They will stay the same
- There is no relationship between bond prices and the interest rate

A 15-year mortgage typically requires higher monthly payments than a 30-year mortgage, but the total interest paid over the life of the loan will be less.
[True, False]

Buying a single company's stock usually provides a safer return than a stock mutual fund.
[True, False]

[Page break]

Information
Would you like to receive free information on how to manage your personal finances and pay off your debt?

- Yes
- No

If you click "Yes", you will receive the information at the end of this survey. If you click "No", you will not receive the information.

[Page break]

How much do you agree or disagree with the statements below?

- There is no excuse for borrowing money
- You should always save up first before buying something
- You can live a good life without borrowing money
- All in all, borrowing money is not worth the cost

[For each item: Strongly agree, Somewhat agree, Neither agree nor disagree, Somewhat disagree, Strongly disagree]

[Page break]

How much do you agree or disagree with the statements below?

- I admire people who own expensive homes, cars, and clothes
- The things I own say a lot about how well I'm doing in life

[For each item: Strongly agree, Somewhat agree, Neither agree nor disagree, Somewhat disagree, Strongly disagree]

[Page break]

In 2019, how much debt did the average American have?
[Slider from $0 to $200,000]

[Page break]

In 2019, what was the share of Americans that had any kind of debt?
[Slider from 0 to 100]

1.F.4.5 Debrief

What do you think was the main hypothesis of this study?
[Text entry field]

If you have any comments related to this study, please write them down in the field below.
[Text entry field]

[Page break]

For your information, you listened to an excerpt from the [Dave Ramsey Show, Modern Mentor Podcast] previously.

[Page break]

Information about personal finances
Here are some suggestions from the Dave Ramsey Show on how to pay off your debt.

[Figure explaining the 7 Baby Steps from the *Dave Ramsey Show*]

Debt Snowball Method
The debt snowball method is a debt-reduction strategy where you pay off debt in order of smallest to largest, gaining momentum as you knock out each remaining balance. When the smallest debt is paid in full, you roll the minimum payment you were making on that debt into the next-smallest debt payment.

- Step 1: List your debts from smallest to largest regardless of interest rate.
- Step 2: Make minimum payments on all your debts except the smallest.
- Step 3: Pay as much as possible on your smallest debt.
- Step 4: Repeat until each debt is paid in full.

Now, before you start arguing about the interest rates, hear us out. If your largest debt has the largest interest rate, it's going to be a long time before you start to see a dent in that crazy balance of yours. But when you stick to the plan (without worrying about interest rates), you're going to be jumping up and down when you pay off that smallest debt super quick. That excitement is what's going to motivate you to keep working hard—all the way to that debt-free finish line.

How useful was this information?
[Very useful, somewhat useful, not useful, not useful at all]

[End of survey]

1.F.4.6 Screenshots

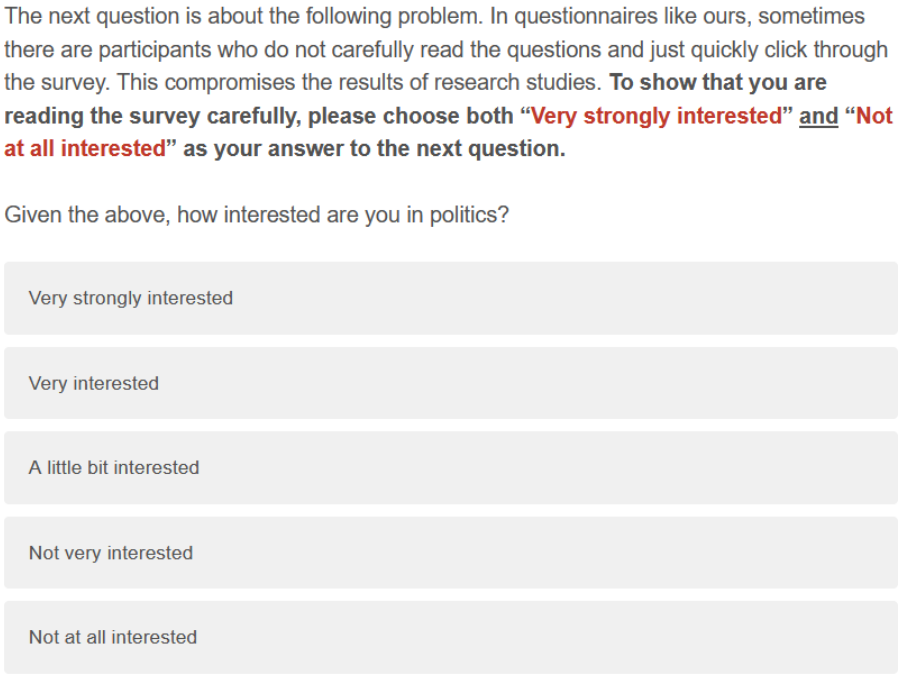

Figure 1.F.2. Attention check

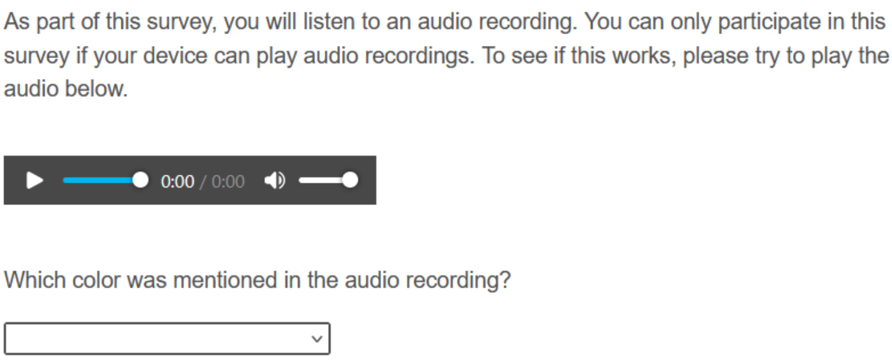

Figure 1.F.3. Audio check

1.F.5 Audio transcript

The control group listened to Episode 277 of the Modern Mentor Podcast by Stever Robbins and published on August 25, 2015. Respondents listened to the 5 minutes and 8 seconds segment from 00:00:09 to 00:05:17. The treatment group listened to an excerpt from the *Dave Ramsey Show*, which was published on March 20, 2017, on the radio show's YouTube channel.[30] Respondents listened to the 5 minute and 4 seconds segment from 00:00:00 to 00:05:04. A verbatim transcript of both excerpts can be found below.

1.F.5.1 Control group

They say you should choose your battles wisely. That makes sense. Consider Napoleon. He chose to fight at Waterloo, and that didn't work out well for him. If he'd chosen more wisely, he might have chosen to fight at Gettysburg. He would have given the Gettysburg Address and had a movie made about him, only instead of starring Daniel-Day Lewis, it would have started Daniel DeVito. One unwisely-chosen battle centuries ago changed the entire course of the Academy Awards centuries later. In our daily lives, choosing battles unwisely means we can waste a lot of time and energy on the wrong thing. This very evening, listener Emily proclaimed on her Facebook wall that she was thrilled that a business celebrity sent her a message. Imagine my surprise to find out she was talking about me! I could have spent time arguing that I'm certainly not a celebrity, and I'm far too humble and modest to deserve such acclaim and adoration. But what would have been the point? I'm sure you'll agree it makes much more sense to accept her statement at face value—as simply a statement of fact—and save my energy for an important battle. Where in your life and work do you fight battles? Why? Are those the right battles? Let's explore how you can make sure you fight less and win more.

I know this sounds obvious, but before going into battle, ask yourself honestly whether you can win. I know you feel you can win but think it through. A coaching client was furious that his biggest customer had stolen some of his technology. He wanted to fight it out in court, but if he won the lawsuit, he'd lose the customer and go out of business. This battle couldn't be won.

It's like trying to get your boyfriend, girlfriend, husband, wife, spousal equivalent, or polyamorous family unit to put the toilet paper roll on with the paper facing the other direction. Not only will you lose that battle, but you'll end up bringing home flowers for a month to repair the damage you made with that foolish, foolish request. You cannot win that battle. So why try?

If you do win, make sure you'll get some benefit from the win. I know people who spend years obsessing over how they were right and Jordan Dinklebert was wrong, but Jordan wouldn't listen and insulted them in front of the entire team. Now they're just waiting for a chance to take revenge. They spend years plotting, and the day they're named employee of the year, halfway through their acceptance speech, they say, "And it's no thanks to Jordan Dinklebert. I was right, you were wrong, and you're really just a big poopie head. So there!" Uh, huh. A poopie head. Well, that little bit of revenge was certainly worth the wait.

Revenge is usually a battle that takes up a lot of resources, and even if you win, you don't really benefit. In Star Trek II: The Wrath of Kahn, Kahn declares, "Revenge is a dish best served cold." Really? Who wants a cold dinner? Revenge is not a dish best served cold. Oreo ice cream cake is a dish best served cold. So what's the lesson here?

30. The full video can be found here: https://www.youtube.com/watch?v=vz-rdaE2uUw

Even if you benefit, make sure you benefit enough to be worth the fight. Take this example: A non-profit organization owned a parcel of undeveloped land. A developer wanted it. He sued the non-profit with a frivolous lawsuit and offered to settle if the non-profit would sell the developer the land for $100,000, which was market price.

The non-profit, on principle, didn't want to give in. But they weren't using the land for anything. And in America, it can cost $20,000 to get a frivolous lawsuit thrown out of court. And the developer, with lawyers on staff, could just sue again. The non-profit realized that even though they could win and keep the land, that win would cost them $20,000. If they didn't fight, they would walk away with $100,000. Were they getting shafted? Yes. But were they smart? Definitely. They chose not to fight a battle that wasn't worth the fight.

Last but not least, consider how else you could spend your time. Even for a battle you can win that is worth the fight, there may be better ways to use your time. One of my clients was spending a lot of time and energy pursuing a contractor who had done shoddy work to his home, defrauding him out of $50,000. When we explored the decision to pursue the case in court, and figured that, given the contractor's resources, my client would recover $25,000 at most, if he won. It would probably take him a day a week for six months, which is 26 days. An entire work month. And that's the best-case scenario.

We looked seriously at all the other opportunities in my client's life and work and realized that he had some business development opportunities that would bring in a six-figure contract if he could work on them full time. The battle with the contractor? He could win. He'd benefit. It would be worth it. But he could spend the same time doing business development instead and make even more money. He chose to forgo the battle and spend his time doing business development. Smart. Next time you start gearing up for a fight, stop. Make sure it's a battle you can win. Make sure you'll benefit if you win it. Make sure the benefit is large, and finally, that there isn't something else you could do instead to get even more benefit elsewhere in your life.

1.F.5.2 Treatment group

If you wanna to win with money, let me give you a good idea. Figure out what most people are doing and run in the other direction. *Run* in the other direction. Most people are broke. Most people look good, and their broke. They spend more than they have coming in. They don't act their wage. They don't live on a plan. They don't agree on spending with their spouse. Their only hope for retirement is that the government, which is well known for its ability to handle money, will take care of them. They don't have money set aside for emergencies. They run credit card debt and student loans and car debt all day, every day. They spend like they're in Congress. Most people are stupid when it comes to money.

70% of Americans are living paycheck to paycheck. The bankruptcy rate is at an all-time high, and foreclosures are rising again. Credit card debt continues to climb, and we have a trillion dollars of student loan debt out there. The average car payment in America today now is 496 dollars over 84 months. That's stupid. Normal in America is broke and stupid. You don't wanna be normal. You wanna be weird. One of the greatest compliments you can get on this show if you call up and I say, "Man, you're weird. I'm looking at weird people. You guys are weird", which means that you're contrary. You are a contrarian. You're perpendicular to the culture. When the culture has lost its way the best thing you can do is be opposite.

Figure out whatever they're doing and do the other thing, right? Because you're not gonna get...you're only going to get what they're getting when you do what they're doing. This is not hard to figure out.

If you keep doing what you've been doing, you're gonna keep getting what you've been getting. You do reap what you sow. You live in a cause-and-effect world, baby. There is no way around this.

So your goal... When I went broke, my goal is to be weird. My goal was to be different. And personal finances is 80% behavior, it's only 20% head knowledge. So, this not some math formula that you have a problem with, this is a person in your mirror. I figured out if I can make the guy in my mirror behave, he can be skinny and rich. He's got issues. And once we realize that behavior is what causes people to handle their money poorly or handle it well, then what we've got to decide is our behaviors. And if you have the same behaviors as broke people have in when it comes to money, you're gonna have the same results as broke people. You're just gonna be another broke person. And some of you are making 250,000 dollars a year and you're broke. You've got no money at all. You've got a mess. Loans coming out your ears. You can't breathe. You run, run, run, run, run, run, run like a rat in a wheel, have a heart attack and die and wonder what happened.

This is no way to live. Buying things you can't afford with money you don't have to impress people you don't really like. Some of you spend an unbelievable amount of money on a car payment to impress someone at a stop light you will never be introduced to. The buddy you felt cool there for about, what, three and a half seconds? Fool.

I've been that fool, that's why I know who he is. I've been that guy, I've been that shallow where I thought that my car actually mattered to somebody. Give me a break. Nobody gives a rip about your car. It, listen, you know what I drive right now? Anything I want. You know why? Because I drove crap for a long time. I drove cars like nobody else would drive. Now I get to drive whatever I wanna drive, and I don't drive them for you. I drive them because I like them. I couldn't give a... care less what you think about what I drive. It's not my problem. It's not your problem either by the way. I'm gonna enjoy. Boy, I like nice cars. But I'm not gonna have a nice car with a stupid car payment on it. It's ridiculous. If your self-esteem is so screwed up that you're doing that then you're gonna struggle with money. You're normal. People spending a bunch of money to act like they're something they're not. What they call in Texas "big hat, no cattle." You need to decide: I don't care what other people think and I'm gonna be weird. Whatever you're doing with money, I'm going to do the opposite thing. And when you decide that, you will start winning with money.

1.F.6 Experimental instructions: Follow-up study

Household Finance Survey 2021

Thank you for your interest in this survey, which is part of a study conducted by researchers from the Bonn Graduate School of Economics. By dedicating **5 minutes** of your time to complete this survey, you help us gain valuable insights about personal finances in America.

Your data will be stored and analyzed in full compliance with the General Data Protection Regulation. In particular, your responses are confidential and no conclusions about your person will be drawn. You can withdraw your consent at any time.

You can read the full privacy policy by clicking here.

Please consent to the processing of your data and our privacy policy.

- I consent
- I do not consent

[Page break]

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This compromises the results of research studies. To show that you are reading the survey carefully, please choose both "Very strongly interested" and "Not at all interested" as your answer to the next question.

Given the above, how interested are you in sports?

- Very strongly interested
- Very interested
- A little bit interested
- Not very interested
- Not at all interested

[Page break]

What is your age?
[Dropdown menu]

What is your gender?

- Male
- Female
- Other / Prefer not to say

In which region do you currently reside?

- Northeast
- Midwest
- South
- West

How many people usually live in your primary residence (including yourself, and excluding non-relatives like roommates or renters)?
[Dropdown menu]

[Page break]

Do you hold any shares of stock in publicly held corporations, stock mutual funds, or investment trusts?

- Yes
- No

How many credit cards do you have?
[Dropdown menu]

[Page break]

We would like to learn more about your primary bank.

[Page break]

How satisfied are you with your primary bank's...?

- Customer service
- Checking account
- Branch and ATM locations

- Mobile banking
- Online banking

[For each item: 5-point scale from “Very satisfied” to “Very dissatisfied”]

[Page break]

How likely are you to recommend your primary bank to a friend or colleague?
[11-point Likert-scale from “Not at all likely” to “Extremely likely”]

[Page break]

Now think about household finances more generally.

[Page break]

How much do you agree or disagree with the statements below?

- There is no excuse for borrowing money
- You should always save up first before buying something
- You can live a good life without borrowing money
- All in all, borrowing money is not worth the cost

[For each item: Strongly agree, Somewhat agree, Neither agree nor disagree, Somewhat disagree, Strongly disagree]

[Page break]

How much do you agree or disagree with the statements below?

- I admire people who own expensive homes, cars, and clothes
- The things I own say a lot about how well I’m doing in life

[For each item: Strongly agree, Somewhat agree, Neither agree nor disagree, Somewhat disagree, Strongly disagree]

Chapter 2

Do People Demand Fact-Checked News? Evidence From U.S. Democrats

Joint with Ingar Haaland and Christopher Roth

Abstract: In a large-scale online experiment with U.S. Democrats, we examine how the demand for a newsletter about an economic relief plan changes when the newsletter content is fact-checked. We first document an overall muted demand for fact-checking when the newsletter features stories from an ideologically aligned source, even though fact-checking increases the perceived accuracy of the newsletter. The average impact of fact-checking masks substantial heterogeneity by ideology: fact-checking reduces demand among Democrats with strong ideological views and increases demand among ideologically moderate Democrats. Furthermore, fact-checking increases demand among all Democrats when the newsletter features stories from an ideologically non-aligned source.

Acknowledgements: We would like to thank the editor of the Journal of Public Economics (Ricardo Perez-Truglia) and three anonymous reviewers for excellent and highly constructive suggestions. We also thank Peter Andre, Joshua Dean, Armin Falk, Lukas Hensel, Yucheng Liang, Matt Lowe, Massimo Morelli, Salvatore Nunnari, Kirill Pogorelskiy, Andrea Prat, Carlo Schwarz, Daniel Sgroi, Andreas Stegmann, Joshua Tasoff, Marc Witte, Florian Zimmermann and seminar participants at the Belief-Based Utility Workshop 2021, Bonn, Warwick, Vienna, and the PERICLES Debate on Political Information at Bocconi for very helpful comments. Financial support from BA/Leverhulme Small Research Grants is gratefully acknowledged. Hrishikesh Iyengar, Apoorv Kanoongo, Melisa Kurtis and Kangli Wang provided excellent research assistance. IRB approval was obtained from the German Association for Experimental Economic Research (GfeW) and the University of Warwick HSSREC (86/20-21). The data collections were pre-registered in the AsPredicted registry. Roth: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866.

2.1 Introduction

Misinformation on mass media is becoming increasingly prevalent (Lazer, Baum, Benkler, Berinsky, Greenhill, et al., 2018). Recent examples of misinformation on mass media include false claims about election fraud in the 2020 U.S. Presidential Election that were widely reported in several mainstream news outlets (Pennycook and Rand, 2021). The rise in misinformation coincides with distrust in the media reaching higher levels than ever, with 56% of Americans saying that the mainstream media is purposely trying to mislead the public with inaccurate reporting.[1] Academics and practitioners alike have suggested fact-checking as one of the main tools to combat misinformation and restore trust in the news (Sell, Hosangadi, Smith, Trotochaud, Vasudevan, et al., 2021). The extent to which fact-checking can be an effective tool to combat misinformation and restore trust in the news crucially depends on the demand for fact-checking services. If consumers—as assumed in many models of media markets—primarily care about the accuracy of the news, news demand should increase when the news content is fact-checked. On the other hand, if consumers also have non-instrumental motives to read news, such as preferences for belief confirmation (Mullainathan and Shleifer, 2005; Di Tella, Perez-Truglia, Babino, and Sigman, 2015; Young, 2016; Faia, Fuster, Pezone, and Zafar, 2021), it is theoretically ambiguous how fact-checking affects the demand for news.

In this paper, we provide evidence on how demand for a newsletter changes when its content is fact-checked. In a large-scale online experiment with more than 4,000 Americans who voted Democratic in the 2020 U.S. Presidential Election, respondents can sign up for a weekly politics newsletter featuring the three top stories about an economic relief plan (the *Biden Rescue Plan*). Whether our respondents sign up for the newsletter is our main outcome of interest. Our key treatment variation is whether respondents are told that all stories featured in the newsletter will be fact-checked. We further cross-randomize whether the newsletter features stories from an ideologically aligned source (*MSNBC*) or a non-aligned news source (*Fox News*). Although focusing exclusively on Democrats limits the generalizability of our results, we made this choice to make sure that the newsletter is equally ideologically aligned for all respondents.

Turning to results, we first establish that our sample of Democrats expects articles featured in the newsletter to contain factual errors and believes that fact-checking increases the accuracy of the newsletter. These results hold irrespective of whether the newsletter features stories from an ideologically aligned or non-aligned source. Our first main result is that demand for a newsletter featuring stories from an ideologically aligned source is largely unaffected by the added fact-checking service: the fact-checking treatment increases newsletter demand by only 1.4 percent-

1. https://www.axios.com/media-trust-crisis-2bf0ec1c-00c0-4901-9069-e26b21c283a9.html (accessed July 9, 2021)

age points. The effect is not statistically significant ($p = 0.382$) and corresponds to a modest 2.7% change in demand compared to the control group mean of 49.7%. It is also relatively precisely estimated given our large sample of more than 4,100 respondents, which gives us an ex-post minimum detectable effect size of 4.4 percentage points (at 80% power). We thus have power to detect relatively modest effect sizes.

Our second main result is that the muted average treatment effect masks substantial heterogeneity by ideology: fact-checking decreases newsletter demand by 6.2 percentage points among Democrats with a strong ideology ($p = 0.021$) and increases demand among moderate Democrats by 4.5 percentage points ($p = 0.018$). These effect sizes correspond to a 10.4% reduction in demand among Democrats with a strong ideology and a 9.9% increase in demand among moderate Democrats (compared to control group means of 59.7% and 45%, respectively), underscoring the economic significance of the effects. Our third main result is that fact-checking increases demand among all Democrats when the newsletter features stories from an ideologically non-aligned source. The treatment increases demand by 10 percentage points on average ($p = 0.016$), which corresponds to a 29.1% increase in demand compared to the control group mean of 34.3%. This underscores the economic significance of the effects.

Our results provide a proof of concept that while fact-checking has the potential to increase the demand for news by increasing its perceived accuracy, it could also have the unintended side effect of reducing the demand for ideologically aligned news among consumers with extreme ideological views, who plausibly have a strong preference for belief confirmation. While these findings could potentially inform the optimal regulation of media markets, one should be careful when trying to generalize from a very specific setting with Democrats only. Our results could plausibly have looked differently if we had run the experiment on a different topic where accuracy concerns are likely to be more important, such as news about COVID-19 vaccine efficacy, or with a sample of Republicans. To draw credible and robust conclusions for policy, future research will need to test the robustness of our findings on the demand for fact-checking across many different settings and samples.

Our paper contributes to several strands of the literature. First, the paper relates to the literature on fact-checking (Barrera, Guriev, Henry, and Zhuravskaya, 2020), debiasing interventions (Cruces, Perez-Truglia, and Tetaz, 2013; Alesina, Carlana, Ferrara, and Pinotti, 2018; Banerjee, Ferrara, and Orozco, 2018; Pennycook and Rand, 2019; Grigorieff, Roth, and Ubfal, 2020; Pennycook, Bear, Collins, and Rand, 2020; Galasso, Morelli, Nannicini, and Stanig, 2021), and misinformation on mass media (Bursztyn, Rao, Roth, and Yanagizawa-Drott, 2020; Pennycook and Rand, 2021). Previous work in this literature has assessed how fact-checking or debiasing interventions affect beliefs and policy views (Nyhan and Reifler, 2010; Nyhan, Porter, Reifler, and Wood, 2019; Barrera et al., 2020; Fehr, Mollerstrom, and Perez-Truglia, 2021; Haaland and Roth, 2021; Haaland, Roth, and Wohlfart, 2021), trust in fact-checking services (Brandtzaeg and Følstad, 2017; Brandtzaeg, Følstad, and Chaparro Domínguez, 2018), and willingness to share false news on social media (Henry, Zhuravskaya, and Guriev, 2021).[2] While these studies have advanced our understanding of how fact-checking affects beliefs and policy views, it is important from a policy perspective to also understand how fact-checking affects people's news demand. We take the first step in this direction by providing evidence on how Democrats' demand for a politics newsletter changes when the newsletter content is fact-checked.

The paper also relates to the literature studying the demand for news (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Prat and Strömberg, 2013; DellaVigna and Ferrara, 2015; Gentzkow, Wong, and Zhang, 2018; Qin, Strömberg, and Wu, 2018). This literature has debated whether people tend to read ideologically aligned news because they have higher trust in ideologically aligned sources or because they want to confirm their existing beliefs (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Druckman and McGrath, 2019). We contribute to this literature by providing a proof of concept that non-instrumental motives, such as preferences for belief confirmation, play a role in driving the demand for ideologically aligned news.

Finally, the paper also relates to the literature on information demand (Zimmermann, 2015; Ganguly and Tasoff, 2016; Falk and Zimmermann, 2017; Golman, Hagmann, and Loewenstein, 2017; Fuster, Perez-Truglia, Wiederholt, and Zafar, 2020; Nielsen, 2020; Tappin, Pennycook, and Rand, 2020; Chopra, Haaland, and Roth, 2021b; Faia et al., 2021; Thaler, 2021).[3] We contribute to this literature by providing evidence on whether Democrats have a preference for more accurate news. Compared to much of the previous literature, our design leverages a more natural outcome, namely people's decision to sign up for a real newsletter covering current political and economic news.

2.2 Sample and experimental design

2.2.1 Sample

We collected the data for the experiment during January and February 2021 in collaboration with Lucid, a data provider commonly used in economic research (Bursztyn, Haaland, Rao, and Roth, 2020; Haaland, Roth, and Wohlfart, 2021). The data was collected in four waves, with about 2,000 respondents per wave and 8,399 re-

2. Work in psychology also studies interventions aiming to reduce the spread of misinformation. For example, attaching warnings to news stories disputed by third-party fact-checkers (Pennycook, Bear, et al., 2020) or using crowdsourcing to generate trust ratings can help consumers identify inaccurate claims (Pennycook and Rand, 2019). While the outcomes considered by this research concern beliefs and trust in news, our focus is on the effects of fact-checking services on the demand for news.

3. See Capozza, Haaland, Roth, and Wohlfart (2021) for a review of the applied literature on information demand.

spondents in total. Each wave was pre-specified in the AsPredicted registry (see Table 2.B.1 for an overview and additional registry information).[4] To make sure that the newsletter was equally ideologically aligned for all respondents, we only recruited respondents who had voted for Joe Biden during the 2020 Presidential Election. Respondents who had voted for another candidate or had not voted at all were immediately screened out of the survey.

One recurring concern about online studies is potentially lower levels of attention among respondents compared to laboratory experiments, which may threaten the internal validity of the study. To address this concern, we included a simple pre-treatment attention check at the beginning of the study (see p. 180 of the *Appendix* for a screenshot). 56% of our respondents passed the attention check, which is very low compared to many other experiments (e.g., 96.4% in Bottan and Perez-Truglia (2020) and 99% in Nathan, Perez-Truglia, and Zentner (2020)). As shown in Section 2.C of the *Appendix*, we also observe much lower data quality among inattentive respondents. We, therefore, focus on attentive respondents in the main specifications, leaving us with a sample of 4,667 respondents.[5,6]

2.2.2 Experimental design

All four waves feature two base treatments that are constant across the waves. In the two base treatments, we vary whether we will fact-check a newsletter featuring the three top stories about the *Biden Rescue Plan* featured on *MSNBC*. On top of this, each wave includes a second set of cross-randomized conditions to assess the robustness of our findings to different variations in the newsletter content and to examine potential mechanisms. Specifically, we vary the framing of the plan (wave 1), the perceived instrumental benefits of the plan (wave 2), whether the newsletter features stories from *MSNBC* or *Fox News* (wave 3), and whether the newsletter features news or opinion pieces (wave 4). Each of the cross-randomized conditions includes a version with fact-checking and one without fact-checking, giving us ten treatments in total across the four waves (with 50% of the respondents being assigned to one of the two base treatments).[7] Section 2.E of the *Appendix* provides screenshots of the full experiment, including all the cross-randomized conditions.

In the experiment, we first measure basic demographics as well as a range of other background characteristics and political views. In the base treatments, respondents are then informed that Congress is debating whether to pass the *Biden Rescue*

4. Each pre-registration was submitted to the AsPredicted registry a few hours before the launch of the respective data collection.

5. Many experimental studies conducted with similar online samples usually screen out inattentive respondents from the outset (e.g., Enke and Graeber, 2019; Haaland and Roth, 2020; Haaland, Roth, and Wohlfart, 2021).

6. We had some minor attrition of 1.1% between the main outcome and the subsequent post-treatment belief measures about newsletter characteristics.

7. Tables 2.B.4–2.B.10 in the *Appendix* assess the integrity of randomization for our treatments.

Plan (the American Rescue Plan Act of 2021) and that the plan has received strong support from liberals but has been criticized by conservatives. We then ask whether they would like to sign up for our weekly newsletter that contains stories about the plan featured on *MSNBC* during the last week.[8] To fix beliefs about the stories featured in the newsletter, we made it clear to respondents that the newsletter would feature "the three top stories about the Biden Rescue Plan featured on MSNBC during the last week." By always focusing on the "three top stories" about the plan, our aim was to make sure that treated respondents did not get the impression that fact-checking affected the selection of articles into the newsletter.

We chose to focus on the *Biden Rescue Plan* because it was heavily featured in the news at the time of the experiment and we believed that demand for stories about the plan would be high. Furthermore, since the *Biden Rescue Plan* included a planned $1,400 stimulus check to all Americans, staying informed about the plan could be instrumentally valuable (e.g. to make optimal saving or investment decisions). We chose to focus on *MSNBC* because it is a well-known liberal outlet that broadly matches the ideological leanings of our respondents. Indeed, in a representative survey of Americans, over 90% who identify *MSNBC* as their primary source of political news are Democrats or lean towards the Democratic party, the highest fraction among any news outlet (Grieco, 2020).

Respondents are randomized into the fact-checking condition (treatment) or the non-fact-checking condition (control). Respondents in the fact-checking condition are informed that "we will fact-check all stories featured in the newsletter and flag those with inaccuracies." Respondents in the non-fact-checking condition are offered the same newsletter but without the fact-checking service.[9] For fact-checking to be valuable, respondents need to have at least some trust in our ability to fact-check the articles. We did not emphasize our affiliation on the decision screen, but the consent form included information about our academic affiliations as "researchers from the University of Bonn, Bergen University, and Warwick University."

Our main outcome of interest is whether people would like to receive our newsletter featuring the three top stories about the *Biden Rescue Plan*. We chose to focus on newsletter subscriptions because newsletters are a popular way of staying informed about politics, with 21% of Americans receiving news from a newsletter over the course of a week (Newman, Fletcher, Schulz, Andi, and Nielsen, 2020). Moreover, by including only the three top articles in our newsletter, we reduce the expected cost of our respondents to stay up to date about the debate of the *Biden*

8. If respondents indicated that they would like to receive our newsletter, we provided them with a link to a website at the end of the survey. The newsletter was published on this website. To accommodate different versions of the newsletter, we created individual websites for each treatment arm (see Figure 2.D.1 for an example). This procedure allowed us to preserve the anonymity of our respondents by circumventing the need to collect email addresses.

9. Figure 2.B.1 of the *Appendix* provides screenshots of the treatment and control condition. Section 2.D provides further details about our fact-checking efforts.

Rescue Plan—both in terms of time costs and search efforts. At the same time, administering the newsletter ourselves allows us to retain sufficient control to vary newsletter characteristics across treatment arms.

We also measure a battery of post-treatment beliefs to assess how fact-checking affected beliefs about different newsletter characteristics, including perceptions of the newsletter's accuracy, the perceived trustworthiness of the newsletter, as well the newsletter's entertainment value, political bias, quality, and complexity. We measure these beliefs using five-point Likert scales. Finally, we elicit beliefs about how many articles featured in the newsletter would contain any factual errors, how many articles they expect to be flagged for inaccuracies, and how much they trust our ability to fact-check the news articles. These questions also allow us to check whether fact-checking affected beliefs about the distribution of articles included in the newsletter.

Discussion of the design. Our base treatments exogenously vary the product characteristics of the newsletter similar to conjoint experiments by offering a fact-checking service to a random subset of respondents. This has a few desirable features. First, by providing additional information about the accuracy of the three top *MSNBC* articles on the *Biden Rescue Plan*, our treatment should not affect beliefs about which articles are featured in the newsletter. We are thus holding beliefs about media bias by omission, filtering, or distortion constant between the treatment group and the control group. Since our treatment should not affect the expected distribution of articles, our design shuts down mechanisms related to rational delegation of costly information acquisition (Suen, 2004; Chan and Suen, 2008). Second, rational agents without non-instrumental motives should prefer fact-checking because they can freely dispose of the additional information. This allows us to rule out prominent mechanisms based on Bayesian updating about the quality of a source that make it difficult to cleanly identify motives with observational data (Gentzkow and Shapiro, 2006). Third, we deliberately offered the fact-checking service ourselves. We truthfully tell our respondents in the treatment group that we will fact-check the newsletter. Our instructions make it clear that we are independent non-partisan researchers.

2.3 Results

2.3.1 Fact-checking of politically aligned news

Descriptives. 49.7% of control group respondents signed up for the newsletter featuring stories from *MSNBC*. The high baseline demand for the newsletter likely reflects that our respondents were interested in staying informed about the outcome of the *Biden Rescue Plan* and saw the newsletter as a convenient tool to receive the most important information. Newsletter demand correlates strongly with the perceived accuracy, entertainment value, quality, and trust in the newsletter (as shown in Figure 2.B.9).

For fact-checking to be valuable in our setting, respondents have to expect at least some factual inaccuracies in the *MSNBC* stories included in the newsletter. Importantly, it is people's subjective expectation of factual inaccuracies—and not the actual prevalence of factual inaccuracies—that determines whether fact-checking should increase the valuation of the newsletter. Figure 2.3.1 uses data from control group respondents to provide descriptive evidence on beliefs about factual inaccuracies in news articles included in the newsletter as well as trust in our ability to fact-check the articles. Figure 2.3.1a shows that 58.8% of the respondents expect at least one article featured in the newsletter with articles from *MSNBC* to contain a factual error. Furthermore, conditional on expecting at least one error, respondents expect 1.6 articles to contain factual errors on average, or slightly more than 50% of all articles.[10]

Another necessary condition for fact-checking to be valuable is that respondents trust our ability to identify potential errors in the articles. As shown in Figure 2.3.1b, we find high levels of trust in our fact-checking ability: 94.9% of the respondents report having at least some trust in our ability to fact-check articles from *MSNBC*, suggesting that our fact-checking treatment has scope to change the perceived accuracy of the newsletter.

Empirical specification. In what follows, we assess how demand for the newsletter changes in response to fact-checking. For that purpose, we estimate the following regression specification using OLS:

$$y_i \;=\; \alpha_0 + \alpha_1 \text{Treatment}_i + \boldsymbol{\alpha}_2 \mathbf{x}_i + \varepsilon_i \tag{2.3.1}$$

where y_i is an indicator taking value one if respondent i signs up for the newsletter and value zero otherwise; Treatment_i is an indicator for whether respondent i is in the fact-checking treatment; $\mathbf{x}_i$ is a vector of control variables[11]; and ε_i is an individual-specific error term. We use robust error terms for inference.

Deviation from the pre-registration. In the main specification, we pool data from all four waves, including the cross-randomized conditions that varied the framing of the plan, the perceived instrumental motives, and whether the newsletter featured

10. We identified factual errors in the articles that were featured in our newsletter. In our main newsletter featuring articles from *MSNBC*, we identified factual errors in two out of 21 articles. The share of articles with an error was 9.5%, which is lower than people's estimate of 30.2%. In our newsletter with *Fox News* articles, 11% of featured articles included an error, which is far below people's expectation of 71.7%. In comparison, Maier (2005) finds an objective error rate of 48% among 4,800 news sources cited in 14 local newspapers.

11. We include the following control variables: gender, education, employment status, log income, Census region, and race and ethnicity. We include wave fixed effects when pooling observations across waves.

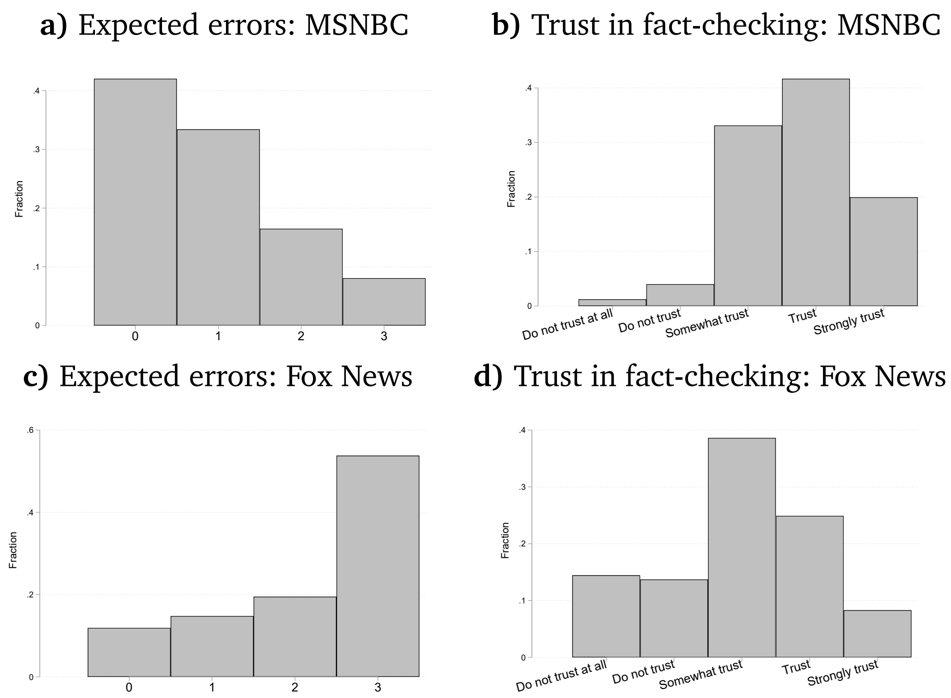

Figure 2.3.1. Expected factual errors and trust in fact-checking

Notes: This figure uses data from control group respondents who passed the attention check. Panel (a) shows the distribution of responses to the question "How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?" Panel (b) shows the distribution of responses to the question "How much do you trust our ability to fact check articles from MSNBC?" Panel (c) and Panel (d) show the corresponding figures for *Fox News*.

news or opinion pieces. These cross-randomized conditions did not differentially affect demand for the newsletter featuring stories from the ideologically aligned source compared to the base treatment (as shown in Table 2.B.14). We deviate from the pre-registration by pooling all results across waves as this allows us to increase the statistical precision of our main estimates and simplify the exposition of our results. A second deviation from the pre-registration is that, motivated by our theoretical model presented in Section 2.A of the Online Appendix, we examine heterogeneity based on the strength of people's ideology. A third deviation from the pre-registration is that, for reasons discussed in Section 2.2.1, we focus on attentive respondents in our main analysis. All pre-registered regressions are reported exactly as pre-specified in Table 2.B.15.

Main effect. Table 2.3.1 presents the main results on how fact-checking affects demand for the newsletter featuring stories from a politically aligned outlet, pooling observations from all waves. Column 1 of Panel A shows the main result of the paper: demand for the newsletter only increases by a non-significant 1.4 percentage points in response to the fact-checking treatment ($p = 0.382$). This effect corresponds to a modest 2.7% change in demand compared to the control group mean of 49.7%. The main effect is relatively precisely estimated given our large sample of more than 4,100 respondents, giving us an ex-post minimum detectable effect size at 80% power of 4.4 percentage points. We thus have power to detect relatively modest effect sizes, suggesting that the average effect of fact-checking on Democrats' demand for news is of relatively low economic importance. Furthermore, as shown in column 2, the muted impact occurs despite a large and statistically significant treatment effect on the perceived accuracy of the newsletter: respondents in the fact-checking condition think that the newsletter has 14.3% of a standard deviation higher accuracy ($p < 0.001$). That treated respondents expect our fact-checking service to increase the overall accuracy of the newsletter is consistent with their high trust in our ability to fact-check the articles (as shown in Figure 2.3.1b). Treated respondents also think that the newsletter has 8.7% of a standard deviation higher trustworthiness ($p = 0.005$). We also see some suggestive evidence that treated respondents associate the newsletter with 4.9% of a standard deviation higher quality ($p = 0.115$) and 5.1% percent of a standard deviation lower left-wing bias ($p = 0.099$), but these effects—while going in the expected direction—are not very large compared to the effect on perceived accuracy. Finally, as shown in columns 6 and 7, it does not seem to be the case that fact-checking affects the perceived complexity ($p = 0.259$) or entertainment value ($p = 0.439$) of the newsletter. Our first main result can be summarized as follows:

Result 1. *On average, people have a muted demand for fact-checking of news from politically aligned sources, despite a significant positive effect of fact-checking on the perceived accuracy of the newsletter.*

Table 2.3.1. Main results: MSNBC

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Main effect							
Treatment	0.014	0.143***	0.087***	0.049	-0.051*	0.035	0.023
	(0.016)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.030)
N	4,109	4,069	4,069	4,069	4,069	4,069	4,069
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.497	0	0	0	0	0	0
Panel B: Strong ideology							
Treatment (a)	-0.062**	0.118**	0.043	0.016	-0.094*	0.027	0.023
	(0.027)	(0.054)	(0.053)	(0.052)	(0.054)	(0.055)	(0.052)
N	1,307	1,299	1,299	1,299	1,299	1,299	1,299
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.597	0	0	0	0	0	0
Panel C: Moderate ideology							
Treatment (b)	0.045**	0.146***	0.097**	0.051	-0.006	0.051	0.010
	(0.019)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.037)
N	2,802	2,770	2,770	2,770	2,770	2,770	2,770
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.450	0	0	0	0	0	0
p-value: a = b	0.001	0.806	0.495	0.638	0.141	0.779	0.808

Notes: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use attentive respondents who were offered a newsletter featuring *MSNBC* articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents with strong ideology (who identify as "very liberal"). Panel C shows results for respondents with moderate ideology (who identify as not "very liberal"). "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining." The outcomes in columns 2–7 are z-scored using the control group mean and standard deviation.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Robustness. We cross-randomized several treatments to assess the robustness of our findings to differences in the content of the newsletter and to examine potential mechanisms. As shown in Table 2.B.14, we find that our main result of a muted demand for fact-checking of ideologically aligned news is robust to varying (i) the framing of the *Biden Rescue Plan* (column 1), (ii) the perceived salience of the financial implications of the plan (column 2), and (iii) the type of articles covered in the newsletter (column 3). Furthermore, as shown in Table 2.B.11, we see very similar point estimates and no significant treatment differences between the base treatments and the pooled cross-randomized treatments. These results suggest that our main finding of an overall muted demand for fact-checking of ideologically aligned news is robust to small variations in the description of the newsletter content.

Heterogeneity by ideology. As discussed in Section 2.A of the *Appendix*, respondents with strong ideological views might assign a larger weight to non-instrumental motives—such as a preference for belief confirmation—than respondents with ideologically moderate views. In this case, we would expect the fact-checking treatment to have an opposite effect on newsletter demand for Democrats with strong and moderate ideological views. To categorize the strength of people's ideological views, we use a pre-treatment question where people report their ideology on a five-point Likert scale from "very liberal" to "very conservative." Throughout the paper, we refer to "very liberal" respondents as those with strong ideological views and to the remaining respondents as moderate respondents.[12] Respondents with strong ideological views hold significantly more extreme policy attitudes than moderate respondents and are, for instance, 54% more likely to "strongly support" the *Biden Rescue Plan*.

Panels B and C of Table 2.3.1 show heterogeneity in treatment effects by ideological views (these effects are also displayed graphically in Panel A of Figure 2.3.2). Panel B of Table 2.3.1 shows treatment effects for respondents with strong ideological views. These respondents significantly *reduce* their demand for the newsletter by 6.2 percentage points in response to the fact-checking treatment ($p = 0.021$, column 1). This corresponds to a 10.4% decline in demand compared to the control group mean of 59.7%, underscoring the economic significance of the effect. The decline in demand occurs even though the respondents perceive the newsletter as 11.8% of a standard deviation more accurate ($p = 0.028$, column 2). These respondents also perceive the fact-checked newsletter as somewhat less left-wing biased ($p = 0.079$, column 5), providing suggestive evidence for a mechanism where respondents with strong ideological views trade off accuracy against non-instrumental

12. 31.8% of our sample rated themselves as "very liberal." Furthermore, consistent with our restriction to focus on respondents who voted for Joe Biden in the 2020 Presidential Election, 93.7% of our respondents rated themselves as either "liberal" or "very liberal." 5.6% rated themselves as "neither liberal nor conservative" and only 0.6% of respondents rated themselves as "conservative" or "very conservative."

utility. Panel C of Table 2.3.1 shows treatment effects for respondents with ideologically moderate views. These respondents significantly increase their demand for the newsletter by 4.5 percentage points in response to the fact-checking treatment ($p = 0.018$, column 1), corresponding to a 9.9% increase in demand compared to a control group mean of 45 percent. Ideologically moderate respondents also perceive the fact-checked newsletter as 14.6% of a standard deviation more accurate ($p < 0.001$, column 2).

Comparing treatment effects in Panel B and Panel C of Table 2.3.1 reveals that we can reject equality of treatment effects on newsletter demand between respondents with strong and moderate ideological views at any conventional level of statistical significance. By contrast, there are no statistically significant differences in treatment effects between the two groups on beliefs about newsletter characteristics, such as accuracy and trust (columns 2–7). Our second main result follows.

Result 2. *Respondents with strong and moderate ideological views respond differently to fact-checking: Despite similar first stage effects on beliefs about newsletter characteristics, respondents with strong ideological views reduce their newsletter demand by 10.4% in response to the fact-checking treatment while ideologically moderate respondents increase their newsletter demand by 9.9%.*

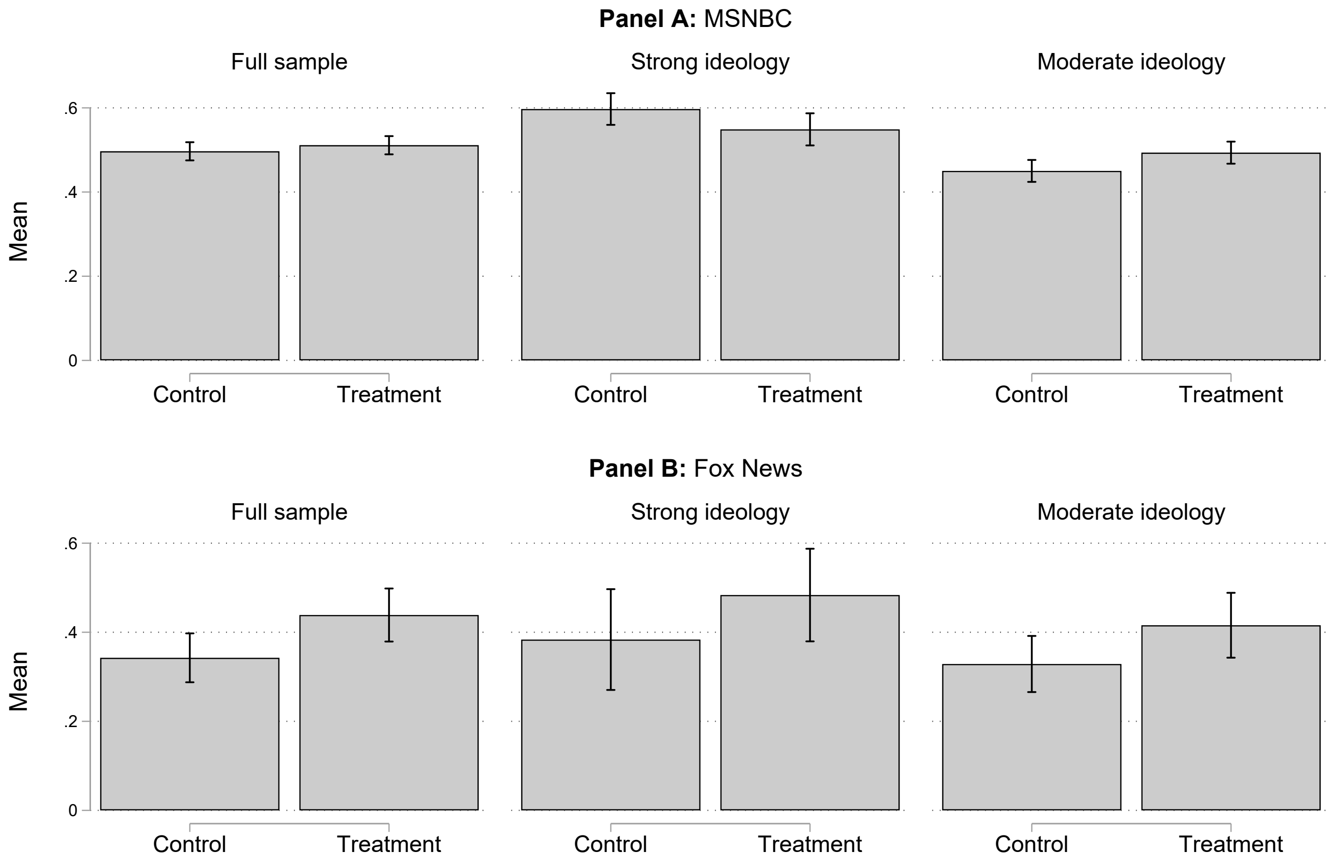

Figure 2.3.2. Treatment effects on demand for the newsletter

Notes: This figure shows newsletter demand (which is a binary variable taking the value one for respondents who said “Yes” to receive the newsletter and zero for those who said “No”) for *MSNBC* (Panel A) and Fox News (Panel B) among attentive respondents. Newsletter demand is shown separately by treatment group for the full sample of Biden voters, respondents with a strong ideology (who identify as “very liberal”), and for respondents with a moderate ideology (who identify as not “very liberal”). 95% confidence intervals are indicated.

2.3.2 Fact-checking of politically non-aligned news

We next study how fact-checking affects demand for a newsletter featuring stories from a politically non-aligned outlet. According to our theoretical framework (Section 2.A of the *Appendix*), fact-checking only creates a trade-off between accuracy and non-instrumental motives when the articles are selected from a politically aligned news outlet. We would therefore expect fact-checking to increase demand for a newsletter featuring stories from a politically non-aligned outlet (Prediction 2 of Section 2.A). To test this prediction, in wave 3, we cross-randomized whether the newsletter featured news articles from *Fox News* instead of *MSNBC* while at the same time holding all other features of the design constant. We chose to focus on *Fox News* because it is a well-known outlet with a conservative leaning. Indeed, in a representative survey of Americans, over 90% who identify *Fox News* as their primary source of political news are Republicans or lean towards the Republican party, the highest fraction among any news outlet (Grieco, 2020).

Descriptives. As expected, we observe a lower demand for news from *Fox News*: 34.3% of control group respondents sign up for the newsletter featuring stories from *Fox News*, compared to 49.7% for *MSNBC*. Given that Biden voters tend to prefer left-wing news, it is reassuring that baseline demand for news from *MSNBC* is 45% higher than for news from *Fox News*. Furthermore, newsletter demand correlates strongly with the perceived accuracy of *Fox News* (as shown in Figure 2.B.10). We next use data from control group respondents to provide descriptive data on beliefs about factual inaccuracies in news articles from *Fox News*. 88.6% of control group respondents expect at least one article to contain factual errors and 53.8% expect every article to contain some errors (Figure 2.3.1c). Furthermore, 73% of the respondents express having at least some trust in our ability to fact-check articles from *Fox News* (Figure 2.3.1d). These descriptives demonstrate a large scope for fact-checking to improve the perceived accuracy of the newsletter.

Main results. Panel A of Table 2.3.2 shows the treatment effects for the 558 respondents in the *Fox News* treatments. Column 1 shows that the fact-checking treatment increases newsletter demand by 10 percentage points ($p = 0.016$). This corresponds to a 29.1% increase in demand relative to the control mean of 34.3%, underscoring the high economic significance of the effect. Respondents in the fact-checking condition also think that the newsletter has 23.1% of a standard deviation higher accuracy ($p = 0.006$, column 2), 15.2% of a standard deviation higher trustworthiness ($p = 0.072$, column 3), and 17.7% of a standard deviation higher quality ($p = 0.038$, column 4).

Table 2.3.2. Main results: Fox News

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Main effect							
Treatment	0.100**	0.231***	0.152*	0.177**	-0.124	-0.076	0.107
	(0.041)	(0.084)	(0.084)	(0.085)	(0.081)	(0.086)	(0.087)
N	558	548	548	548	548	548	548
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.343	0	0	0	0	0	0
Panel B: Strong ideology							
Treatment (a)	0.064	0.195	0.117	0.227	-0.208	-0.035	0.265
	(0.079)	(0.158)	(0.163)	(0.172)	(0.151)	(0.159)	(0.176)
N	164	163	163	163	163	163	163
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.384	0	0	0	0	0	0
Panel C: Moderate ideology							
Treatment (b)	0.095*	0.224**	0.141	0.147	-0.062	-0.081	0.022
	(0.049)	(0.101)	(0.099)	(0.097)	(0.097)	(0.102)	(0.096)
N	394	385	385	385	385	385	385
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.329	0	0	0	0	0	0
p-value: a = b	0.732	0.953	0.973	0.637	0.381	0.826	0.202

Notes: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use attentive respondents who were offered a newsletter featuring *Fox News* articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents with strong ideology (who identify as "very liberal"). Panel C shows results for respondents with moderate ideology (who identify as not "very liberal"). "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining." The outcomes in columns 2–7 are z-scored using the control group mean and standard deviation.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Heterogeneity by ideology. Table 2.3.2 presents treatment effects for Democrats with strong ideology (Panel B) and Democrats with moderate ideology (Panel C). While focusing on these subsamples substantially reduces our power to detect statistically significant effects, especially for respondents with strong ideology, we find broadly similar patterns for both groups. As shown in column 1, treated respondents with strong and moderate ideology increase their demand for the newsletter by 6.4 percentage points ($p = 0.42$) and 9.5 percentage points ($p = 0.056$), respectively (these results are also shown graphically in Panel B of Figure 2.3.2). The increase in demand among both groups is consistent with the theory that the trade-off between instrumental and non-instrumental motives disappears when the newsletter features stories from a politically non-aligned source. Furthermore, as shown in columns 2–7, treatment effects on beliefs about newsletter characteristics, including perceived accuracy, are also similar in magnitude and with no significant differences between the two groups. This leads to our third main result:

Result 3. *All respondents, irrespective of their ideological leanings, increase their demand for the newsletter from a politically non-aligned source in response to the fact-checking treatment.*

2.3.3 Alternative mechanisms

In this section, we discuss a series of mechanisms, which might be operating in this setting, but which are unlikely to explain the patterns in our data.

Confidence and ideology. Empirically, we find that both respondents with moderate and strong ideology expect a more accurate newsletter if it is fact-checked (column 2 of Table 2.3.1). However, respondents with strong ideology, who hold strong prior beliefs about the world, might be very confident that they can detect any inaccuracies in reporting themselves. While overconfidence might decrease the perceived added-value of fact-checking services, it cannot strictly decrease the valuation of the newsletter. This would require an additional feature such as a large cost of processing information.

Updating about source quality. People might update about the quality of the underlying source of the newsletter when they learn that the source is fact-checked. For instance, people could think that fact-checking implies that the underlying source is of low quality (hence the need for a fact-check). To address this potential concern, we elicited expected errors from the underlying source of the newsletter. If anything, we actually see that our respondents in the fact-check condition expect fewer errors from the underlying source (Table 2.B.16).

Cognitive constraints. Furthermore, since fact-checking in our context does not affect the selection of articles in the newsletter, we can—to the extent that fact-checking itself is not perceived as cognitively costly—change beliefs about accuracy while holding cognitive costs constant. Even if our respondents perceive fact-checking as cognitively costly (which we consider unlikely as column 6 of Table 2.3.1 shows that fact-checking does not affect the perceived complexity of the newsletter), the heterogeneity by the strength of people's ideological views as well as the heterogeneity by the ideological leanings of the outlet suggest that cognitive constraints are not driving the observed patterns in our data.

Demand effects. While the high baseline demand for the newsletter featuring stories from *MSNBC* to some degree could reflect experimenter demand effects, this is not an issue for estimating treatment effects unless there is differential experimenter demand across treatment and control. While the between-design should not make it salient that we are interested in how fact-checking affects newsletter demand, we cannot rule out that some respondents nonetheless realized that we were studying fact-checking and adjusted their behavior accordingly. However, recent evidence suggests that demand effects are not a major concern in online experiments (Quidt, Haushofer, and Roth, 2018).

2.3.4 Expert survey

Lastly, we wanted to examine how experts expect the demand for the newsletter to change in response to fact-checking of the newsletter content. The results from this study can potentially inform a policy maker's trade-off between following expert advice on fact-checking in a different setting and conducting new experiments (DellaVigna and Pope, 2018). For this purpose, we conducted a survey in March 2021 among leading academic researchers in the areas of media and behavioral economics. We compiled a list of 93 experts who attended major conferences in economics.[13] Our final sample consists of 65 experts, corresponding to a response rate of 70 percent.[14] After providing the expert participants with information about the sample, design, and experimental instructions (including screenshots of the key treatment screens), we elicit their predictions about the effect of fact-checking on the demand for news for *MSNBC* and *Fox News*. For both outlets, we inform experts about baseline demand for the newsletter among respondents in the control group and then elicit their beliefs about newsletter demand among respondents in the treatment group.

13. These conferences include the briq Workshop on Beliefs, the NBER Summer Institute in Political Economy, and the Stanford Institute for Theoretical Economics (SITE) Summer Workshop (Experimental Economics and Psychology & Economics sessions).

14. 25% of these experts are Full Professor, 15% are Associate Professor, and 34% are Assistant Professors, 14% are postdoctoral researchers, and only 12% of respondents in our sample are PhD students. Among non-respondents, 65.5% are Full Professors, 14% are Associate Professors, 18% are Assistant Professors, and 4% are PhD students. This suggests lower response rates among full professors compared to assistant professors, PhD students, and postdoctoral researchers.

Figure 2.B.11 of the *Appendix* shows the results from the expert survey. As shown in Figure 2.B.11a, we observe a wide dispersion in expert beliefs about the impact of fact-checking on the demand for news with a mean absolute deviation of seven percentage points between expert opinions and actual treatment effects. The heterogeneity in expert beliefs suggests that there is substantial expert disagreement about the relative importance of different motives to read the news, such as the importance of accuracy motives versus belief utility motives. As shown in Figure 2.B.11b, expert beliefs on average closely resemble the actual treatment effects. As in DellaVigna and Pope (2018), our findings demonstrate a strong wisdom-of-crowds effect: while there is substantial disagreement within the expert sample, experts on average correctly predict the effects of fact-checking on newsletter demand.

2.4 Concluding remarks

This paper studies how fact-checking affects the demand for news. The main result of the paper is that Democrats have a muted demand for fact-checking of a newsletter featuring ideologically aligned news, even though fact-checking increases the perceived accuracy of the newsletter. This average effect masks substantial heterogeneity: Fact-checking decreases demand for politically aligned news among Democrats with strong ideological views and increases demand among ideologically moderate Democrats. Furthermore, fact-checking increases the demand for a newsletter with politically non-aligned news for all Democrats irrespective of the strength of their ideological leanings.

Our findings provide a proof of concept that non-instrumental motives play a role in driving the demand for ideologically aligned news. These findings have relevance for theories of media markets. In particular, our findings are inconsistent with theories in which all consumers primarily care about the accuracy of the news and point to the relevance of theories incorporating non-instrumental motives, such as a preference for belief confirmation. Furthermore, while one should be careful not to overgeneralize from a very specific setting, our findings suggest that fact-checking services can have very heterogeneous effects on the demand for news. While our study provides the first step to understand how fact-checking affects the demand for news, our results could be specific to our chosen sample and setting. Future research should study how fact-checking affects the demand for news across a range of different settings and samples to generate useful lessons for policy-makers.

References

Alesina, Alberto, Michela Carlana, Eliana La Ferrara, and Paolo Pinotti. 2018. "Revealing stereotypes: Evidence from immigrants in schools." Working paper. National Bureau of Economic Research. [119]

Banerjee, Abhijit, Eliana Ferrara, and Victor Orozco. 2018. "The entertaining way to behavioral change: Fighting HIV with MTV." Working paper. Working paper. [119]

Barrera, Oscar, Sergei Guriev, Emeric Henry, and Ekaterina Zhuravskaya. 2020. "Facts, alternative facts, and fact checking in times of post-truth politics." *Journal of Public Economics* 182: 104123. [119]

Bottan, Nicolas L., and Ricardo Perez-Truglia. 2020. "Choosing Your Pond: Location Choices and Relative Income." *Review of Economics and Statistics*, (10): 1–46. [121]

Brandtzaeg, Petter Bae, and Asbjørn Følstad. 2017. "Trust and distrust in online fact-checking services." *Communications of the ACM* 60 (9): 65–71. [119]

Brandtzaeg, Petter Bae, Asbjørn Følstad, and Maria Ángeles Chaparro Domínguez. 2018. "How Journalists and Social Media Users Perceive Online Fact-Checking and Verification Services." *Journalism Practice* 12 (9): 1109–29. [119]

Bursztyn, Leonardo, Ingar Haaland, Aakaash Rao, and Christopher Roth. 2020. "Disguising Prejudice: Popular Rationales as Excuses for Intolerant Expression." *University of Chicago, Becker Friedman Institute for Economics Working Paper, 2020-73*, [120]

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." *SSRN Electronic Journal*, DOI: 10.2139/ssrn.3580487. [119]

Capozza, Francesco, Ingar Haaland, Christopher Roth, and Johannes Wohlfart. 2021. "Studying Information Acquisition in the Field: A Practical Guide and Review." Discussion paper 124. ECONtribute. [120]

Chan, Jimmy, and Wing Suen. 2008. "A Spatial Theory of News Consumption and Electoral Competition." *Review of Economic Studies* 75 (3): 699–728. [123]

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2019. "Do People Value More Informative News?" *SSRN Working Paper 3342595*, DOI: http://dx.doi.org/10.2139/ssrn.3342595.

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2021a. "Do People Demand Fact-Checked News? Evidence From U.S. Democrats." *ECONtribute Discussion Paper No. 121*,

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2021b. "Do People Value More Informative News?" Working Paper. [120]

Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment." *Journal of Public Economics* 98: 100–112. [119]

DellaVigna, Stefano, and Eliana La Ferrara. 2015. "Economic and Social Impacts of the Media." In *Handbook of Media Economics*. Edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg. Vol. 1B, Handbook of Media Economics. North-Holland, 723–68. [120]

DellaVigna, Stefano, and Devin Pope. 2018. "Predicting Experimental Results: Who Knows What?" *Journal of Political Economy* 126 (6): 2410–56. DOI: 10.1086/699976. [134, 135]

Di Tella, Rafael, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman. 2015. "Conveniently Upset: Avoiding Altruism by Distorting Beliefs About Others' Altruism." *American Economic Review* 105 (11): 3416–42. [118]

Druckman, James N., and Mary C. McGrath. 2019. "The evidence for motivated reasoning in climate change preference formation." *Nature Climate Change* 9 (2): 111–19. [120]

Enke, Benjamin, and Thomas Graeber. 2019. "Cognitive Uncertainty." NBER Working Papers 26518. National Bureau of Economic Research. URL: https://ideas.repec.org/p/nbr/nberwo/26518.html. [121]

Faia, Ester, Andreas Fuster, Vincenzo Pezone, and Basit Zafar. 2021. "Biases in information selection and processing: Survey evidence from the pandemic." Working paper. National Bureau of Economic Research. [118, 120]

Falk, Armin, and Florian Zimmermann. 2017. "Beliefs and Utility: Experimental Evidence on Preferences for Information." *Working Paper*, [120]

Fehr, Dietmar, Johanna Mollerstrom, and Ricardo Perez-Truglia. 2021. "Your Place in the World: The demand for National and Global Redistribution." *American Economic Journal: Economic Policy* (forthcoming), [119]

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2020. "Expectations with Endogenous Information Acquisition: An Experimental Investigation." *Review of Economics and Statistics* (forthcoming), [120]

Galasso, Vincenzo, Massimo Morelli, Tommaso Nannicini, and Piero Stanig. 2021. "Fighting Populism on Its Own Turf: Experimental Evidence." *Working Paper*, [119]

Ganguly, Ananda, and Joshua Tasoff. 2016. "Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future." *Management Science* 63 (12): 4037–60. [120]

Gentzkow, Matthew A., and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy* 114 (2): 280–316. [120, 123]

Gentzkow, Matthew A., Michael B. Wong, and Allen T. Zhang. 2018. "Ideological Bias and Trust in Information Sources." *Working Paper*, [120]

Golman, Russell, David Hagmann, and George Loewenstein. 2017. "Information Avoidance." *Journal of Economic Literature* 55 (1): 96–135. [120]

Grieco, Elizabeth. 2020. "Americans' main sources for political news vary by party and age." https://pewrsr.ch/3dHYbmB (accessed 02-24-2021). [122, 131]

Grigorieff, Alexis, Christopher Roth, and Diego Ubfal. 2020. "Does Information Change Attitudes Toward Immigrants?" *Demography* 57 (3): 1–27. [119]

Haaland, Ingar, and Christopher Roth. 2020. "Labor market concerns and support for immigration." *Journal of Public Economics* 191: 104256. DOI: https://doi.org/10.1016/j.jpubeco.2020.104256. [121]

Haaland, Ingar, and Christopher Roth. 2021. "Beliefs About Racial Discrimination and Support for Pro-Black Policies." *Review of Economics and Statistics* (forthcoming), [119]

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2021. "Designing Information Provision Experiments." *Journal of Economic Literature* (forthcoming), [119–121]

Henry, Emeric, Ekaterina Zhuravskaya, and Sergei Guriev. 2021. "Checking and sharing alt-facts." *American Economic Journal: Economic Policy* (forthcoming), [120]

Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. "The science of fake news." *Science* 359 (6380): 1094–96. [118]

Maier, Scott R. 2005. "Accuracy Matters: A Cross-Market Assessment of Newspaper Error and Credibility." *Journalism & Mass Communication Quarterly* 82 (3): 533–51. [124]

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–53. [118, 120, 141]

Nathan, Brad C., Ricardo Perez-Truglia, and Alejandro Zentner. 2020. "My Taxes are Too Darn High: Why Do Households Protest Their Taxes?" Working paper. National Bureau of Economic Research. [121]

Newman, Nic, Richard Fletcher, Anne Schulz, Simge Andi, and Rasmus Kleis Nielsen. 2020. "Reuters Institute Digital News Report 2020." Working paper. Reuters Institute for the Study of Journalism. [122]

Nielsen, Kirby. 2020. "Preferences for the resolution of uncertainty and the timing of information." *Journal of Economic Theory* 189: 105090. [120]

Nyhan, Brendan, Ethan Porter, Jason Reifler, and Thomas J. Wood. 2019. "Taking fact-checks literally but not seriously? The effects of journalistic fact-checking on factual beliefs and candidate favorability." *Political Behavior* 42: (3), 939–60. [119]

Nyhan, Brendan, and Jason Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions." *Political Behavior* 42: 303–0. [119]

Pennycook, Gordon, Adam Bear, Evan T. Collins, and David G. Rand. 2020. "The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings." *Management Science* 66 (11): 4944–57. [119, 120]

Pennycook, Gordon, and David G. Rand. 2019. "Fighting misinformation on social media using crowdsourced judgments of news source quality." *Proceedings of the National Academy of Sciences* 116 (7): 2521–26. [119, 120]

Pennycook, Gordon, and David G. Rand. 2021. "Research note: Examining false beliefs about voter fraud in the wake of the 2020 Presidential Election." *Harvard Kennedy School (HKS) Misinformation Review*, [118, 119]

Prat, Andrea, and David Strömberg. 2013. "The Political Economy of Mass Media." In *Advances in Economics and Econometrics: Applied Economics*. Edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel. Vol. 2, Advances in Economics and Econometrics. Cambridge University Press, Cambridge, 135–87. [120]

Qin, Bei, David Strömberg, and Yanhui Wu. 2018. "Media Bias in China." *American Economic Review* 108 (9): 2442–76. [120]

Quidt, Jonathan de, Johannes Haushofer, and Christopher Roth. 2018. "Measuring and Bounding Experimenter Demand." *American Economic Review* 108 (11): 3266–302. [134]

Sell, Tara Kirk, Divya Hosangadi, Elizabeth Smith, Marc Trotochaud, Prarthana Vasudevan, Gigi Kwik Gronvall, Yonaira Rivera, Jeannette Sutton, Alex Ruiz, and Anita Cicero. 2021. "National Priorities to Combat Misinformation and Disinformation for COVID-19 and Future Public Health Threats: A Call for a National Strategy." Baltimore, MD: Johns Hopkins Center for Health Security. [118]

Suen, Wing. 2004. "The Self-Perpetuation of Biased Beliefs." *Economic Journal* 114 (495): 377–96. [123]

Tappin, Ben M., Gordon Pennycook, and David G. Rand. 2020. "Thinking clearly about causal inferences of politically motivated reasoning: why paradigmatic study designs often undermine causal inference." *Current Opinion in Behavioral Sciences* 34: 81–87. [120]

Thaler, Michael. 2021. "The "Fake News" Effect: An Experiment on Motivated Reasoning and Trust in News." *Working Paper*, [120]

Young, Eric. 2016. "A new understanding: What makes people trust and rely on news." American Press Institute. [118]

Zimmermann, Florian. 2015. "Clumped or Piecewise? Evidence on Preferences for Information." *Management Science* 61 (4): 740–53. [120]

Summary of the Appendices

Section 2.A presents our theoretical framework.

Section 2.B contains additional tables and figures. Table 2.B.1 provides an overview of the four experimental waves. Table 2.B.2 provides summary statistics for demographic variables for the attentive sample and separately for each wave. Table 2.B.3, Table 2.B.4, Table 2.B.5, Table 2.B.6, Table 2.B.7, Table 2.B.8, Table 2.B.9, and Table 2.B.10 provide balance tests for our treatment manipulations based on observables. Table 2.B.11 shows the main treatment effects on *MSNBC* newsletter demand for the full attentive sample as well as separately for respondents in the base treatments and in the extra treatments. Table 2.B.12 shows the main treatment effects on *MSNBC* newsletter demand separately for respondents with strong ideology and moderate ideology using the full sample of respondents (including respondents who did not pass the attention check). Table 2.B.13 shows the main treatment effects on *Fox News* newsletter demand separately for respondents with strong ideology and moderate ideology using the full sample of respondents (including inattentive respondents). Table 2.B.14 shows interaction effects between our base treatment and our additional treatments separately for attentive respondents and inattentive respondents. Table 2.B.15, which includes all our pre-registered regressions, shows interaction effects between our base treatment and our additional treatments by each wave using the full sample (including inattentive respondents). Table 2.B.16 shows treatment effects on expected errors. Table 2.B.17 shows differences in covariates between respondents who signed up for the newsletter and those who did not. Figure 2.B.1 provides a screenshot of the key treatment screens. Figure 2.B.2 shows the distribution of beliefs about factual errors and trust in our ability to fact-check the articles included in the newsletter for the full sample (including inattentive respondents). Figure 2.B.3 shows treatment effects graphically using the full sample (including inattentive respondents). Figure 2.B.4 shows the distribution of beliefs about factual errors and trust in our ability to fact-check news articles by respondent's ideology and the news outlet. Figure 2.B.5 shows the distribution of beliefs about different newsletter characteristics by ideology and news outlet. Figure 2.B.6 shows the evolution of demand for our newsletter over time. Figure 2.B.7 shows the results from simultaneously interacting our main treatment with respondent ideology and a vector of controls for the *MSNBC* newsletter. Figure 2.B.8 shows the results from simultaneously interacting our main treatment with respondent ideology and a vector of controls for the *Fox News* newsletter. Figure 2.B.9 and Figure 2.B.10 show correlates of the demand for news from *MSNBC* and *Fox News*, respectively. Figure 2.B.11 shows results from the expert survey.

Section 2.C compares the sample of attentive and inattentive respondents. Table 2.C.1 provides summary statistics separately for attentive and inattentive respondents. Table 2.C.2 shows the main treatment effects on *MSNBC* newsletter demand for the full sample as well as separately for attentive and inattentive respondents.

Table 2.C.3 shows the main treatment effects on *Fox News* newsletter demand for the full sample as well as separately for attentive and inattentive respondents. Figure 2.C.1 shows correlations between newsletter demand and beliefs about newsletter characteristics separately for attentive and inattentive respondents.
Section 2.D provides further details about the newsletter and our fact-checking efforts, including an example of how our newsletter looked like.
Section 2.E provides screenshots of the experimental instructions.

Appendix 2.A Theoretical framework

This section lays out a simple Bayesian model of news consumption where agents face a trade-off between instrumental and non-instrumental concerns. Based on this framework, we generate predictions for how fact-checking could affect the demand for news. There is an unobserved binary state $\theta \in \{L, R\}$ that captures the desirability of a policy proposed by Democrats, which in our experiment is the *Biden Rescue Plan*. The agent, a Biden voter, has a prior belief $q \geq 1/2$ that the plan will have positive overall consequences, i.e., $\theta = L$.

Politically aligned outlet. The agent can read a politically biased newsletter that contains a binary news article $n \in \{L, R\}$. We start with the case of a newsletter featuring articles from a politically aligned news outlet. The agent expects this outlet to always report L if indeed $\theta = L$. However, with probability p, the agent thinks the newsletter will report L even if $\theta = R$. Thus, p captures the perceived left-wing bias in reporting.[15]

The agent has to take a binary action $a \in \{L, R\}$ with incentives to match the state. A relevant action could be how much to save, which depends on the expected stimulus check from the *Biden Rescue Plan*. Specifically, she receives utility α if her action matches the state.[16] Without reading the newsletter, the agent will always choose L given her prior belief, which generates expected utility of αq. Now, reading the newsletter increases the matching probability by $(1-q)(1-p)$. The newsletter's instrumental value, u_I, is therefore

$$u_I \;=\; \alpha(1-q)(1-p). \tag{2.A.1}$$

The agent may also receive non-instrumental utility from reading politically aligned news. For example, the agent might have a preference for news that confirms her prior beliefs about the world (Mullainathan and Shleifer, 2005), which might conflict with her preference for more accurate news. In our model, the agent receives utility β from reading news articles that confirm her prior belief that $\theta = L$. Given her beliefs, the expected non-instrumental utility is then

$$u_B \;=\; \beta\left(q + (1-q)p\right). \tag{2.A.2}$$

15. The agent's belief about biased reporting—not the actual probability of distortion—determines the anticipated utility consequences of reading the newsletter. This allows us to also capture cases where respondents have biased beliefs. Moreover, by continuity, our results also hold if $P(n = L \mid \theta = L) = \tau$ for large τ.

16. An alternative interpretation is that the agent intrinsically cares about learning the truth. Then α captures the intrinsic value from holding accurate beliefs about the world.

Now suppose the newsletter is fact-checked by an external party. The fact-checker will flag all inaccurate articles, thereby identifying the share of articles p that is left-wing biased.[17] This has two opposing effects. On the one hand, the instrumental utility increases by $\alpha(1-q)p$ because the newsletter now fully reveals the state. On the other hand, the non-instrumental utility from biased reporting decreases by $\beta(1-q)p$, implying a net change of the agent's valuation of the newsletter by

$$\Delta u_{\text{aligned}} = (\alpha - \beta)(1-q)p. \tag{2.A.3}$$

This generates the following prediction:

Prediction 1. Fact-checking a newsletter featuring articles from a politically aligned news outlet will, (i), decrease the demand for news among respondents with stronger non-instrumental motives ($\alpha < \beta$) and, (ii), increase the demand for news among respondents with stronger instrumental motives ($\alpha > \beta$).

For example, people with strong ideological views might care more about the non-instrumental utility from belief confirmation than people with moderate views. In this case, we would expect fact-checking to have a polarizing effect on demand.

Politically non-aligned outlet. We finally consider the case of a politically non-aligned news outlet. Here, the agent expects the news outlet to report R if $\theta = R$ and to report R with probability p' if $\theta = L$. Thus, p' captures the perceived right-wing biased of the news outlet. First, suppose the agent decides to read the outlet's newsletter. In this case, we can derive her posterior belief $\hat{q}(n)$ that $\theta = L$ from Bayes' rule:

$$\hat{q}(n) = \begin{cases} 1 & \text{if } n = L \\ \frac{qp'}{1-q+p'q} & \text{if } n = R \end{cases} \tag{2.A.4}$$

The agent will find it optimal to choose $a = R$ after reading $n = R$ only if $\hat{q}(R) \leq \frac{1}{2}$, which is the case if $(1+p')q \leq 1$. Thus, after reading the article n, it is optimal to choose $a = n$ if $(1+p')q \leq 1$, and $a = L$ otherwise.

Again, fact-checking will increase the instrumental value of the newsletter by identifying the share of articles p' that incorrectly reports about the state θ. However, fact-checking now increases the non-instrumental utility as well because factual inaccuracies consist of reporting R although $n = L$ would have been correct. In total, the agent's valuation of the newsletter changes by

$$\Delta u_{\text{opposed}} = \big(\alpha + \beta\big)qp' + \alpha \max\{0, 1-(1+p')q\}. \tag{2.A.5}$$

due to the fact-checking, which implies:

17. We obtain qualitatively similar results if fact-checking is only able to flag inaccuracies with probability τ. Moreover, the results also hold if fact-checking only decreases the non-instrumental utility from inaccurate reports to $\beta' < \beta$.

Prediction 2. Fact-checking a newsletter featuring articles from a politically non-aligned news outlet will increase the demand for news.

Proof. The proof is by case distinction. First, consider the case where $(1+p')q \leq 1$. In this case, the agent's action will match the state whenever $n = \theta$, which happens with probability $1-q+q(1-p')$. Relative to always choosing $a = L$, the newsletter provides instrumental utility of $u_I = \alpha(1-q+q(1-p'))$, and non-instrumental utility of $\beta q(1-p')$. Now, fact-checking will increase the instrumental value by $\alpha qp'$ and the non-instrumental utility by $\beta qp'$. In total, the agent's valuation increases by $\Delta u = (\alpha+\beta)qp'$. Second, consider the case where $(1+p')q > 1$. In this case, the agent will always choose L. Thus, the instrumental value of the newsletter is $u_I = 0$. Thus, while the effect of fact-checking on the non-instrumental utility is identical to the previous case, fact-checking will now increase the instrumental value by $\alpha(1-q)$ because it is now optimal to choose $a = n$. Thus, the total change in the agent's valuation is

$$\Delta u \;=\; \alpha(1-q) + \beta qp' \;=\; (\alpha+\beta) + \alpha(1-(1+p')q). \tag{2.A.6}$$

Thus, we have shown that for politically non-aligned outlets, the effect of fact-checking on the agent's valuation of a newsletter is positive and given by

$$\Delta u_{\text{opposed}} \;=\; (\alpha+\beta) + \alpha \max\{0, 1-(1+p')q\}, \tag{2.A.7}$$

which is strictly positive. This concludes the proof. □

Appendix 2.B Additional Tables and Figures

Table 2.B.1. Overview of experimental waves

Wave	Sample	Date	Extra treatments	Pre-analysis plan
Wave 1	*n* = 2,086	Jan 21–22	Non-polarized topic	aspredicted.org/blind.php?x=vk4ap3
Wave 2	*n* = 2,097	Jan 22–26	Instrumental value	aspredicted.org/blind.php?x=j22u5z
Wave 3	*n* = 2,054	Feb 15–16	Right-wing outlet	aspredicted.org/blind.php?x=qe6ad3
Wave 4	*n* = 2,162	Feb 16–18	Commentary	aspredicted.org/blind.php?x=zs5ht9

Notes: This table provides an overview of the four experimental waves. All four waves feature the two base treatments (demand for *Biden Rescue Plan* with or without fact-check). In addition, each wave has an extra set of treatments.

Table 2.B.2. Summary statistics: Attentive respondents

	(1) Full sample	(2) Wave 1	(3) Wave 2	(4) Wave 3	(5) Wave 4
Male	0.400	0.380	0.442	0.426	0.349
Age	43.267	41.016	44.128	47.360	40.574
White	0.765	0.749	0.781	0.783	0.746
Log income	10.834	10.838	10.848	10.825	10.823
College education	0.864	0.860	0.867	0.879	0.850
Full-time employee	0.447	0.457	0.460	0.385	0.487
Northeast	0.229	0.234	0.205	0.239	0.239
Midwest	0.225	0.219	0.225	0.257	0.200
West	0.221	0.217	0.249	0.193	0.227
South	0.324	0.330	0.321	0.312	0.334
Hispanic	0.102	0.113	0.096	0.091	0.108
Observations	4,667	1,322	1,183	1,146	1,016

Notes: This table displays the mean value of basic covariates for the full attentive sample (column 1) and separately for each wave (columns 2–5). “Male” is a binary variable with value one for male respondents. “Age” is age of the respondent. “White” is a binary variable with value one if the respondent selected “Caucasian/White.” “Log income” is coded continuously as the log of the income bracket’s midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). “College education” is a binary variable taking value one if the respondent selected “Some college, no degree,” “Associates degree,” “Bachelor’s degree,” or “Post-graduate degree.” “Full-time employee” is a binary variable taking value one if the respondent is a full-time employee. “Northeast,” “Midwest,” “West” and “South” are binary variables with value one if the respondent lives in the respective region. “Hispanic” is a binary variable with value one if the respondent is Hispanic.

Table 2.B.3. Test of balance for attentive respondents: Treatment vs. control

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Male	0.40	0.40	0.663	4667
Age	43.31	43.22	0.851	4667
Log of income	10.85	10.82	0.168	4667
South	0.32	0.33	0.227	4667
West	0.22	0.22	0.587	4667
Northeast	0.24	0.22	0.024	4667
White	0.76	0.77	0.392	4667
College	0.87	0.86	0.576	4667
Full-time employee	0.45	0.44	0.703	4667
Hispanic	0.10	0.11	0.383	4667

Notes: This table provides a balance test for the fact-checking treatment using attentive respondents from all waves. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

Table 2.B.4. Test of balance for full sample: Treatment vs. control

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Male	0.44	0.44	0.639	8399
Age	40.08	39.98	0.790	8399
Log of income	10.75	10.76	0.858	8399
South	0.33	0.35	0.182	8399
West	0.20	0.20	0.857	8399
Northeast	0.24	0.23	0.065	8399
White	0.65	0.67	0.231	8399
College	0.81	0.81	0.946	8399
Full-time employee	0.48	0.48	0.973	8399
Hispanic	0.15	0.16	0.662	8399

Notes: This table provides a balance test for the fact-checking treatment using observations from all waves. “Male” is a binary variable with value one for male respondents. “Age” is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). “Log of income” is coded continuously as the logarithm of the income bracket’s midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). “South,” “West,” and “Northeast” are binary variables with value one if the respondent lives in the respective region. “White” is a binary variable with value one if the respondent selected “Caucasian/White.” “College education” is a binary variable taking value one if the respondent selected “Some college, no degree,” “Associates degree,” “Bachelor’s degree,” or “Post-graduate degree.” “Full-time employee” is a binary variable taking value one if the respondent is a full-time employee. “Hispanic” is a binary variable with value one if the respondent is Hispanic.

Table 2.B.5. Test of balance for attentive respondents with a strong ideology: Treatment vs. control

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Male	0.43	0.38	0.054	1471
Age	40.54	40.26	0.737	1471
Log of income	10.84	10.80	0.405	1471
South	0.34	0.31	0.152	1471
West	0.22	0.21	0.621	1471
Northeast	0.25	0.25	0.822	1471
White	0.76	0.78	0.438	1471
College	0.87	0.87	0.849	1471
Full-time employee	0.50	0.47	0.302	1471
Hispanic	0.12	0.11	0.528	1471

Notes: This table provides a balance test for the fact-checking treatment using attentive respondents with a strong ideology from all waves. “Male” is a binary variable with value one for male respondents. “Age” is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). “Log of income” is coded continuously as the logarithm of the income bracket’s midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). “South,” “West,” and “Northeast” are binary variables with value one if the respondent lives in the respective region. “White” is a binary variable with value one if the respondent selected “Caucasian/White.” “College education” is a binary variable taking value one if the respondent selected “Some college, no degree,” “Associates degree,” “Bachelor’s degree,” or “Post-graduate degree.” “Full-time employee” is a binary variable taking value one if the respondent is a full-time employee. “Hispanic” is a binary variable with value one if the respondent is Hispanic.

Table 2.B.6. Test of balance for attentive respondents with a moderate ideology: Treatment vs. control

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Male	0.39	0.40	0.428	3196
Age	44.63	44.54	0.894	3196
Log of income	10.86	10.83	0.265	3196
South	0.30	0.34	0.015	3196
West	0.22	0.22	0.744	3196
Northeast	0.24	0.20	0.004	3196
White	0.76	0.77	0.608	3196
College	0.87	0.86	0.428	3196
Full-time employee	0.43	0.43	0.764	3196
Hispanic	0.09	0.10	0.111	3196

Notes: This table provides a balance test for the fact-checking treatment using attentive respondents with a moderate ideology from all waves. “Male” is a binary variable with value one for male respondents. “Age” is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). “Log of income” is coded continuously as the logarithm of the income bracket’s midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). “South,” “West,” and “Northeast” are binary variables with value one if the respondent lives in the respective region. “White” is a binary variable with value one if the respondent selected “Caucasian/White.” “College education” is a binary variable taking value one if the respondent selected “Some college, no degree,” “Associates degree,” “Bachelor’s degree,” or “Post-graduate degree.” “Full-time employee” is a binary variable taking value one if the respondent is a full-time employee. “Hispanic” is a binary variable with value one if the respondent is Hispanic.

Table 2.B.7. Test of balance: Neutral versus polarized framing

	Neutral (a)	Polarized (b)	P-value(a - b)	Observations
Male	0.37	0.39	0.468	1322
Age	41.78	40.29	0.105	1322
Log of income	10.86	10.81	0.237	1322
South	0.32	0.34	0.432	1322
West	0.21	0.22	0.502	1322
Northeast	0.25	0.22	0.315	1322
White	0.75	0.75	0.897	1322
College	0.86	0.86	0.664	1322
Full-time employee	0.43	0.48	0.105	1322
Hispanic	0.10	0.12	0.321	1322

Notes: This table provides a balance test for neutral and polarized framing of the policy proposal using attentive respondents from wave 1. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

Table 2.B.8. Test of balance: High instrumental value versus neutral framing

	Instrumental (a)	Neutral (b)	P-value(a - b)	Observations
Male	0.42	0.47	0.122	1183
Age	44.75	43.45	0.186	1183
Log of income	10.84	10.86	0.763	1183
South	0.32	0.32	0.851	1183
West	0.25	0.24	0.745	1183
Northeast	0.21	0.20	0.559	1183
White	0.80	0.76	0.148	1183
College	0.86	0.87	0.494	1183
Full-time employee	0.44	0.49	0.076	1183
Hispanic	0.10	0.09	0.629	1183

Notes: This table provides a balance test for instrumental value treatment vs neutral (base) treatment using attentive respondents from wave 2. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

Table 2.B.9. Test of balance: *Fox News* versus *MSNBC*

	Fox News (a)	MSNBC (b)	P-value(a - b)	Observations
Male	0.45	0.40	0.139	1146
Age	47.12	47.59	0.657	1146
Log of income	10.83	10.82	0.893	1146
South	0.32	0.30	0.432	1146
West	0.19	0.20	0.696	1146
Northeast	0.24	0.24	0.935	1146
White	0.79	0.78	0.543	1146
College	0.88	0.88	0.812	1146
Full-time employee	0.40	0.37	0.261	1146
Hispanic	0.10	0.08	0.371	1146

Notes: This table provides a balance test for the *Fox News* versus *MSNBC* treatment using attentive respondents from wave 3.riates "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

Table 2.B.10. Test of balance: Opinion versus news

	Opinion (a)	News (b)	P-value(a - b)	Observations
Male	0.35	0.35	0.802	1016
Age	40.27	40.86	0.557	1016
Log of income	10.80	10.84	0.397	1016
South	0.34	0.33	0.807	1016
West	0.23	0.23	0.984	1016
Northeast	0.23	0.25	0.404	1016
White	0.75	0.74	0.570	1016
College	0.84	0.86	0.344	1016
Full-time employee	0.47	0.51	0.223	1016
Hispanic	0.08	0.13	0.013	1016

Notes: This table provides a balance test for the opinion versus news section variation using attentive respondents from wave 4. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

Table 2.B.11. Heterogeneity by base vs. extra treatments: MSNBC

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Main effect							
Treatment	0.014	0.143***	0.087***	0.049	-0.051*	0.035	0.023
	(0.016)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.030)
N	4,109	4,069	4,069	4,069	4,069	4,069	4,069
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.497	0	0	0	0	0	0
Panel B: Base treatments							
Treatment (a)	0.029	0.135***	0.084**	0.044	-0.102**	0.044	0.068*
	(0.021)	(0.041)	(0.040)	(0.041)	(0.040)	(0.041)	(0.040)
N	2,354	2,336	2,336	2,336	2,336	2,336	2,336
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.491	0	0	0	0	0	0
Panel C: Extra treatments							
Treatment (b)	-0.005	0.152***	0.090*	0.051	0.014	0.008	-0.033
	(0.024)	(0.048)	(0.048)	(0.048)	(0.049)	(0.049)	(0.046)
N	1,755	1,733	1,733	1,733	1,733	1,733	1,733
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.504	0	0	0	0	0	0
p-value: a = b	0.291	0.769	0.953	0.911	0.062	0.562	0.099

Notes: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use respondents that were offered a newsletter featuring *MSNBC* articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents assigned to the base treatments. Panel C shows pooled from the cross-randomized conditions in wave 1 (different framing of the plan), wave 2 (higher perceived instrumental motives of the plan), and wave 4 (opinion stories about the plan). "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased.." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining."

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.B.12. Heterogeneous treatment effects between respondents with strong and moderate views: MSNBC (full sample)

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Strong ideology							
Treatment (a)	-0.024	0.077**	-0.012	-0.014	-0.041	-0.040	-0.008
	(0.018)	(0.038)	(0.039)	(0.038)	(0.037)	(0.038)	(0.037)
N	2,592	2,571	2,571	2,571	2,571	2,571	2,571
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.657	0	0	0	0	0	0
Panel B: Moderate ideology							
Treatment (b)	0.019	0.068**	0.037	0.043	-0.010	0.044	0.008
	(0.014)	(0.030)	(0.030)	(0.029)	(0.029)	(0.029)	(0.028)
N	4,779	4,723	4,723	4,723	4,723	4,723	4,723
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.500	0	0	0	0	0	0
p-value: a = b	0.062	0.745	0.320	0.254	0.435	0.089	0.731

Notes: This table uses data from all respondents (including inattentive ones) and shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use respondents that were offered a newsletter featuring *MSNBC* articles. Panel A shows results for respondents with strong ideology (who identify as “very liberal”) and Panel B shows results for respondents with moderate ideology (who identify as not “very liberal”). “Treatment” is a binary variable taking value one if the articles in the newsletter are fact-checked. “News demand” is a binary variable taking the value one for respondents who said “Yes” to receive the newsletter and zero for those who said “No.” “Accuracy” of the newsletter is measured on a 5-point scale from “Very inaccurate” to “Very accurate.” “Trust” is the trustworthiness of the newsletter and measured on a 5-point scale from “Not trustworthy at all” to “Very trustworthy.” “Quality” of the newsletter is measured on a 5-point scale from “Very low quality” to “Very high quality.” “Left-wing bias” is measured on a 5-point scale from “Very right-wing biased” to “Very left-wing biased..” “Complexity” of the newsletter articles is measured on a 5-point scale from “Very simple” to “Very complex.” “Entertainment” of the newsletter is measured on a 5-point scale from “Not entertaining at all” to “Very entertaining.”
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.B.13. Heterogeneous treatment effects between respondents with strong and moderate views: Fox News (full sample)

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Strong ideology							
Treatment (a)	0.099*	0.167	0.182*	0.241**	-0.144	-0.124	0.173*
	(0.052)	(0.104)	(0.103)	(0.107)	(0.104)	(0.104)	(0.104)
N	329	328	328	328	328	328	328
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.548	0	0	0	0	0	0
Panel B: Moderate ideology							
Treatment (b)	0.062*	0.157**	0.151**	0.146**	-0.127*	-0.089	0.120
	(0.037)	(0.074)	(0.074)	(0.074)	(0.073)	(0.078)	(0.074)
N	699	682	682	682	682	682	682
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.402	0	0	0	0	0	0

Notes: This table uses data from all respondents (including inattentive ones) and shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use respondents that were offered a newsletter featuring *Fox News* articles. All regressions use respondents that were offered a newsletter featuring *MSNBC* articles. Panel A shows results for respondents with strong ideology (who identify as "very liberal") and Panel B shows results for respondents with moderate ideology (who identify as not "very liberal"). "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased.."
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.B.14. Interaction analysis: Base versus extra treatments

	Interactant:			
	(1) Neutral frame	(2) Instrumental value frame	(3) Opinion piece	(4) Right-wing outlet
Panel A: Attentive respondents				
Treatment	0.026	0.027*	0.029*	0.014
	(0.016)	(0.016)	(0.015)	(0.016)
Treatment × Interactant	-0.017	-0.028	-0.060	0.081*
	(0.042)	(0.043)	(0.048)	(0.044)
Interactant	0.020	-0.016	0.029	-0.145***
	(0.035)	(0.036)	(0.038)	(0.035)
N	4,667	4,667	4,667	4,667
Controls	Yes	Yes	Yes	Yes
Control group mean	0.491	0.491	0.491	0.491
Panel B: Inattentive				
Treatment	-0.006	-0.003	0.001	-0.012
	(0.016)	(0.016)	(0.017)	(0.017)
Treatment × Interactant	0.009	-0.012	-0.039	0.057
	(0.050)	(0.049)	(0.042)	(0.046)
Interactant	-0.033	-0.003	0.081**	-0.083**
	(0.042)	(0.039)	(0.036)	(0.039)
N	3,732	3,732	3,732	3,732
Controls	Yes	Yes	Yes	Yes
Control group mean	0.625	0.625	0.625	0.625

Notes: This table shows OLS regression where the dependent variable is demand for the newsletter (taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No"). We pool respondents across waves. "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked (base treatment). In each column, we interact the base treatment with a different additional treatment. The interactants are binary variables taking value one if a respondent was assigned to the condition of the additional treatment that differed from the base experiment. In each column, we include indicator variables for the additional treatments that are not explored in the interaction analysis.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.B.15. Interaction of the base treatment and the additional treatments

	Interactant:			
	(1) Neutral frame	(2) Instrumental value frame	(3) Opinion piece	(4) Right-wing outlet
Treatment (a)	-0.024	0.011	0.019	0.032
	(0.031)	(0.030)	(0.030)	(0.030)
Treatment × Interactant (b)	0.032	-0.023	-0.039	0.043
	(0.043)	(0.043)	(0.042)	(0.043)
Interactant	-0.016	-0.018	0.056*	-0.102***
	(0.030)	(0.030)	(0.030)	(0.030)
N	2,086	2,097	2,162	2,054
Controls	Yes	Yes	Yes	Yes
Control group mean	0.552	0.570	0.549	0.532
P-value: a + b = 0	0.783	0.702	0.481	0.013

Notes: This table shows OLS regression where the dependent variable is demand for the newsletter (taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No"). Each column uses only observations from that particular wave, i.e, column *k* uses respondents from wave *k*. "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked (base treatment). In each column, we interact the base treatment with the additional treatment in that particular wave. The interactants are binary variables taking value one if a respondent was assigned to the condition of the additional treatment that differed from the base experiment.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.B.16. Treatment effect on expected errors

	Attentive respondents		Full sample	
	(1)	(2)	(3)	(4)
	MSNBC	Fox News	MSNBC	Fox News
Treatment	-0.120***	-0.264***	-0.066***	-0.127*
	(0.029)	(0.097)	(0.023)	(0.070)
N	4,039	539	7,236	996
Z-scored	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Control group mean	0.906	2.152	1.072	1.900

Notes: This table shows OLS regression where the dependent variable are the respondent's expectation about the number of articles that contain factual inaccuracies in reporting, which can range from 0 to 3. "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. Columns 1 and 2 show results for attentive respondents, while columns 3 and 4 show results for the full sample of Biden voters (including inattentive respondents). Columns 1 and 3 use respondents that were offered a newsletter featuring *MSNBC* articles, while columns 2 and 4 those that were offered a newsletter featuring *Fox News* articles.

* p < 0.10, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses.

Table 2.B.17. Selection

	Newsletter demand	No newsletter demand	P-value	Observations
Male	0.42	0.36	0.003	2059
Age	43.68	41.63	0.006	2059
Log of income	10.84	10.81	0.469	2059
South	0.33	0.33	0.768	2059
West	0.22	0.23	0.594	2059
Northeast	0.24	0.19	0.014	2059
White	0.74	0.79	0.008	2059
College	0.87	0.86	0.556	2059
Full-time employee	0.47	0.43	0.070	2059
Hispanic	0.13	0.09	0.008	2059

Notes: This table shows the characteristics of respondents who signed up for the newsletter ("Newsletter demand") and those who did not ("No newsletter demand") among attentive control group respondents who were offered the newsletter featuring articles from *MSNBC*. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log of income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South," "West," and "Northeast" are binary variables with value one if the respondent lives in the respective region. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Hispanic" is a binary variable with value one if the respondent is Hispanic.

a) Newsletter: Control group

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in January and February 2021.

Would you like to receive our newsletter?

- ○ Yes
- ○ No

b) Newsletter: Treatment group

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in January and February 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

- ○ Yes
- ○ No

Figure 2.B.1. Experimental instructions: Newsletter about the *Biden Rescue Plan*

Notes: These figures provide the experimental instructions used to describe the politics newsletter to respondents in the control group (Panel A) and in the treatment group (Panel B) for the case of a politically aligned outlet. The original instructions did not include the red highlighting in Panel B. For the politically non-aligned outlet, we replaced MSNBC with Fox News.

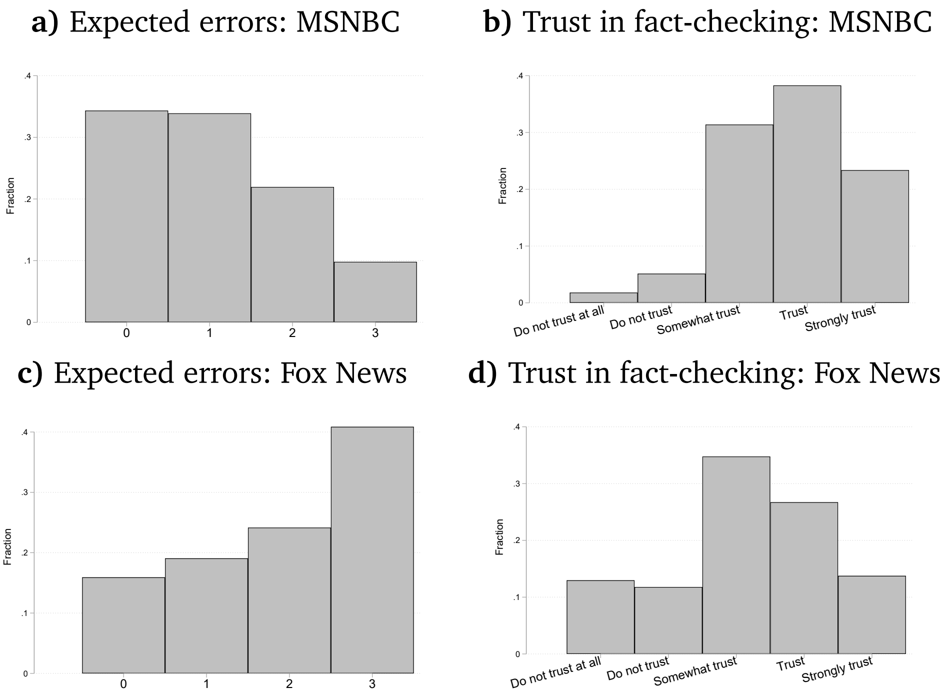

Figure 2.B.2. Expected factual errors and trust in fact-checking: Full sample

Notes: This figure uses data from control group respondents (including those who did not pass the attention check). Panel 2.B.2a shows the distribution of responses to the question "How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?" Panel 2.B.2b shows the distribution of responses to the question "How much do you trust our ability to fact check articles from MSNBC?" Panel 2.B.2c and Panel 2.B.2d show the corresponding figures for *Fox News*.

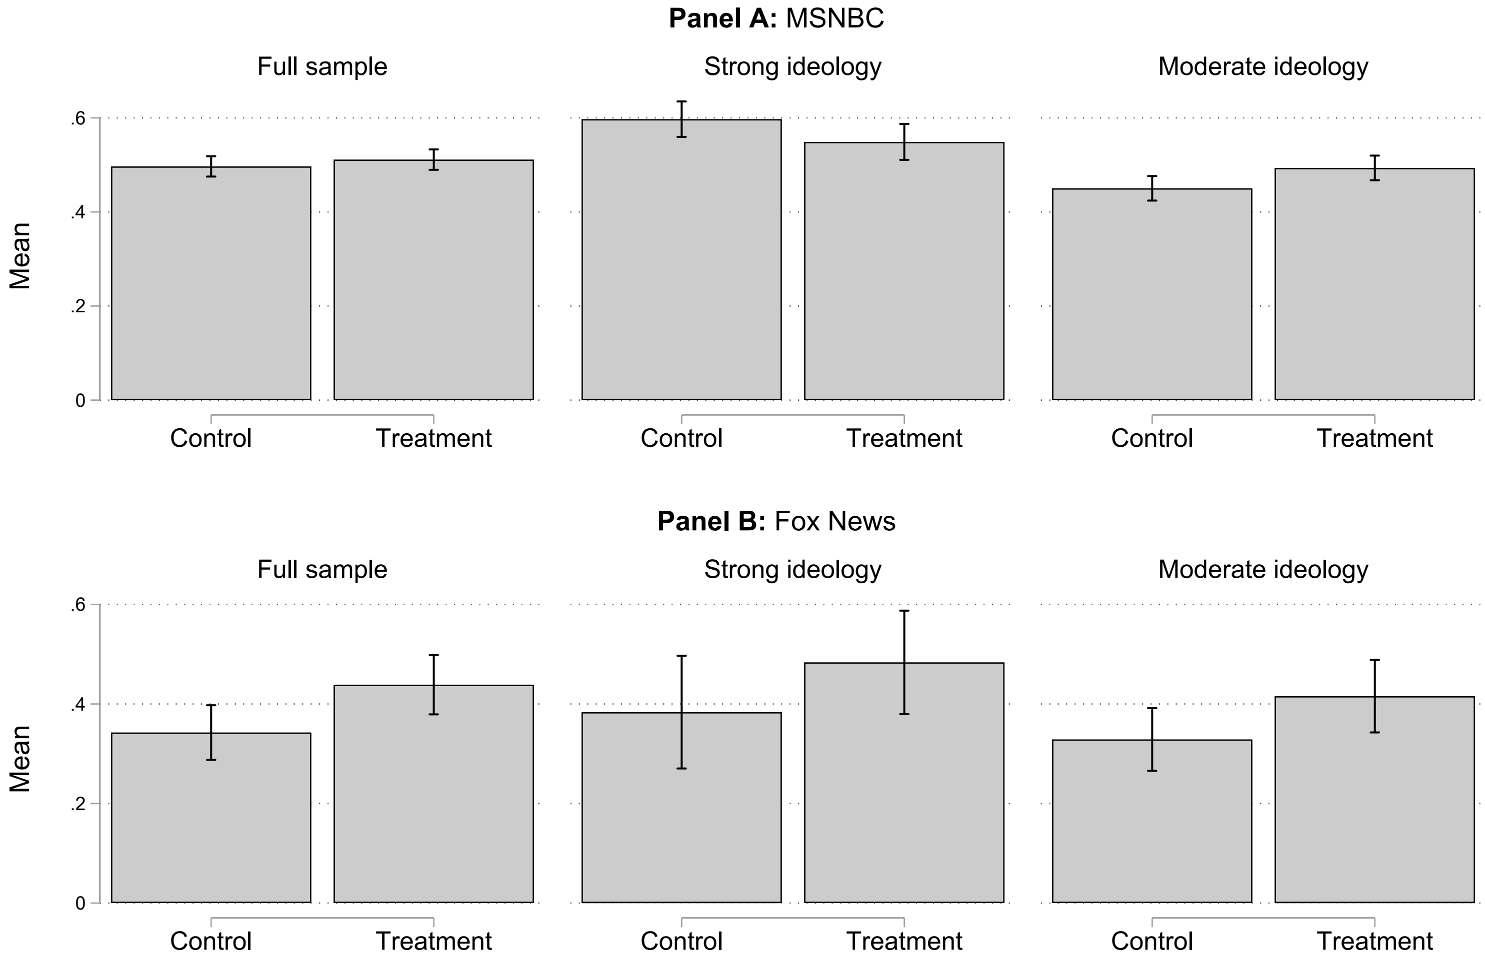

Figure 2.B.3. Treatment effects on demand for the newsletter: Full sample

Notes: This figure shows newsletter demand for *MSNBC* (Panel A) and Fox News (Panel B) using all respondents (including inattentive respondents). Newsletter demand is shown separately by treatment group for the full sample of Biden voters, respondents with a strong ideology (who identify as “very liberal”), and for respondents with a moderate ideology (who identify as not “very liberal”). 95% confidence intervals are indicated.

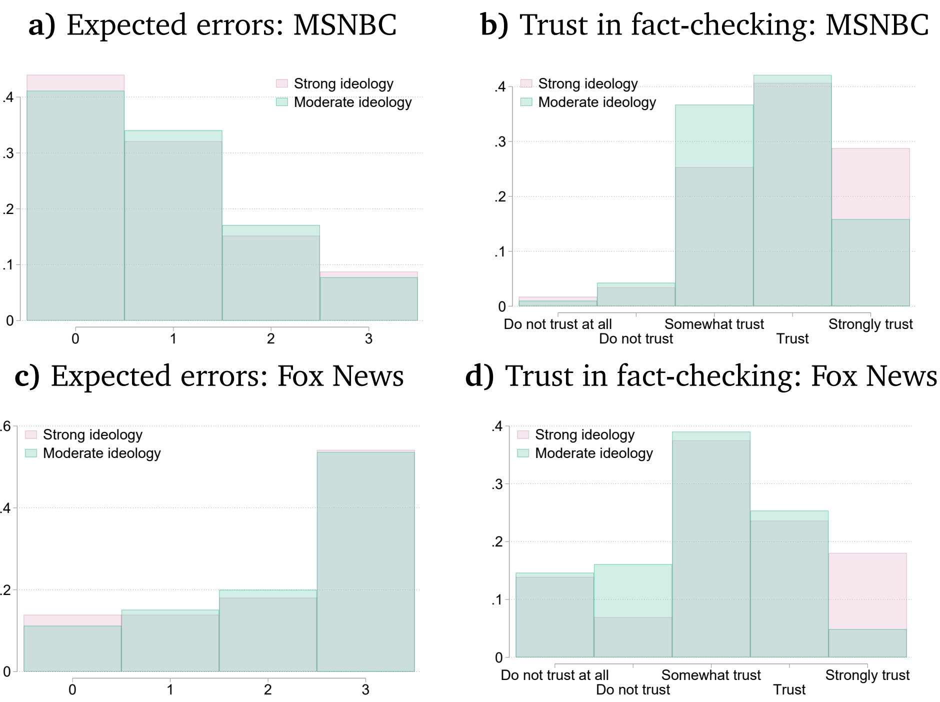

Figure 2.B.4. Expected factual errors and trust in fact-checking ability separately by ideology

Notes: This figure uses data from control group respondents who passed the attention check. Panel 2.B.4a shows the distribution of responses to the question "How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?" Panel 2.B.4b shows the distribution of responses to the question "How much do you trust our ability to fact check articles from MSNBC?" Panel 2.B.4c and Panel 2.B.4d show the corresponding figures for Fox News.

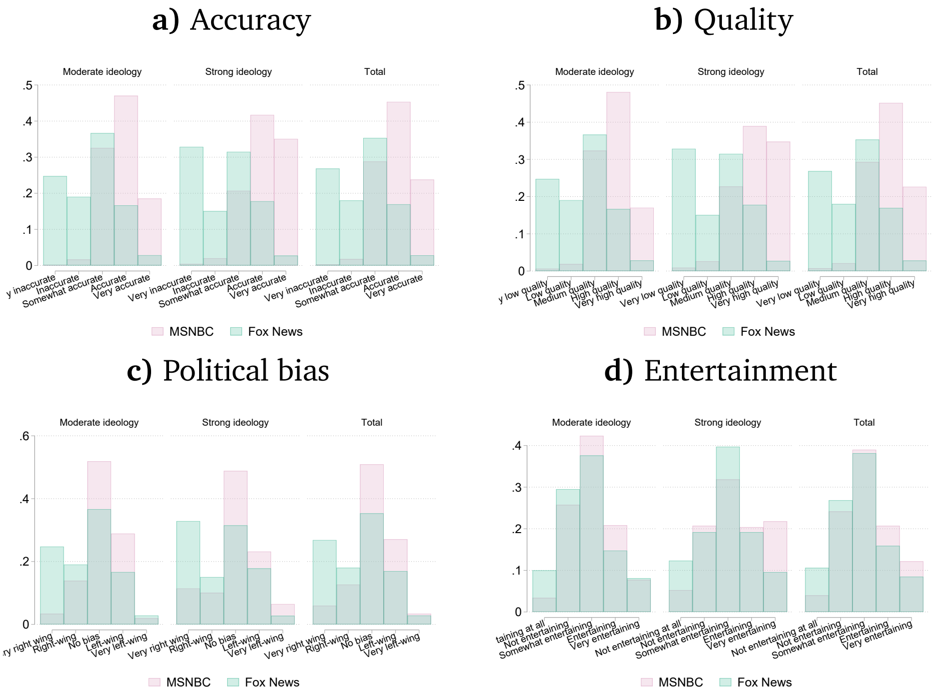

Figure 2.B.5. Beliefs about newsletter characteristics

Notes: This figure uses data from control group respondents who passed the attention check. Figure 2.B.5a shows the distribution of responses to the question "How accurate do you expect the newsletter to be?." Figure 2.B.5b shows the distribution of responses to the question "What quality would you expect the newsletter to have?." Figure 2.B.5c shows the distribution of responses to the question "What kind of political bias do you expect the newsletter to have?." Figure 2.B.5d shows the distribution of responses to the question "How entertaining do you expect the newsletter to be?." Each panel separately shows the distribution of responses for respondents with a strong ideology, moderate ideology and the full sample.

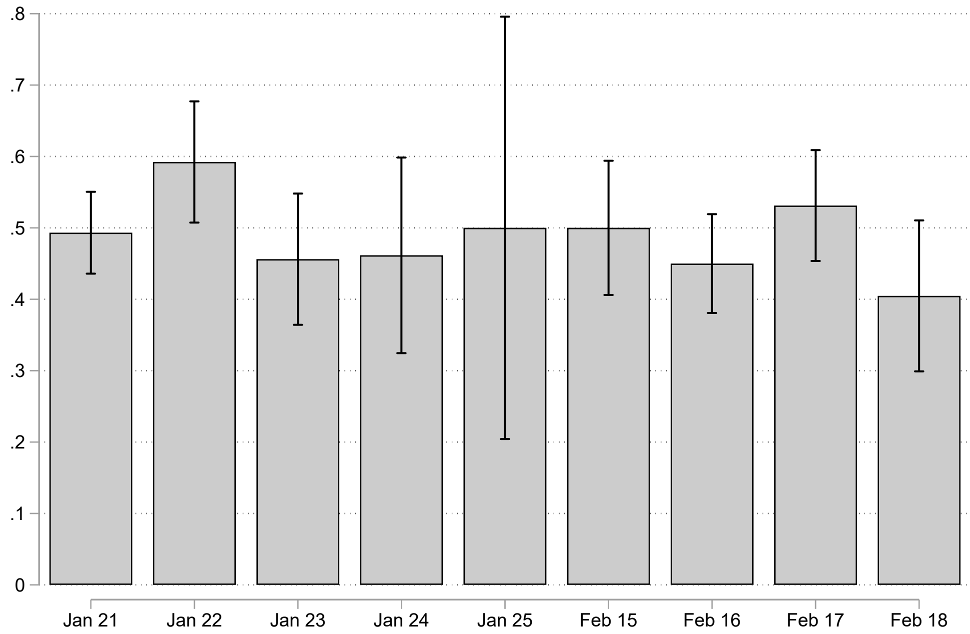

Figure 2.B.6. Newsletter demand over time

Notes: This figure uses data from control group respondents in the base treatment who passed the attention check. The vertical bars indicate the fraction of respondents who signed up for the newsletter. 95% confidence intervals are indicated. The date indicators are not jointly significantly different from zero in a regression with newsletter demand as the dependent variable (p = 0.191).

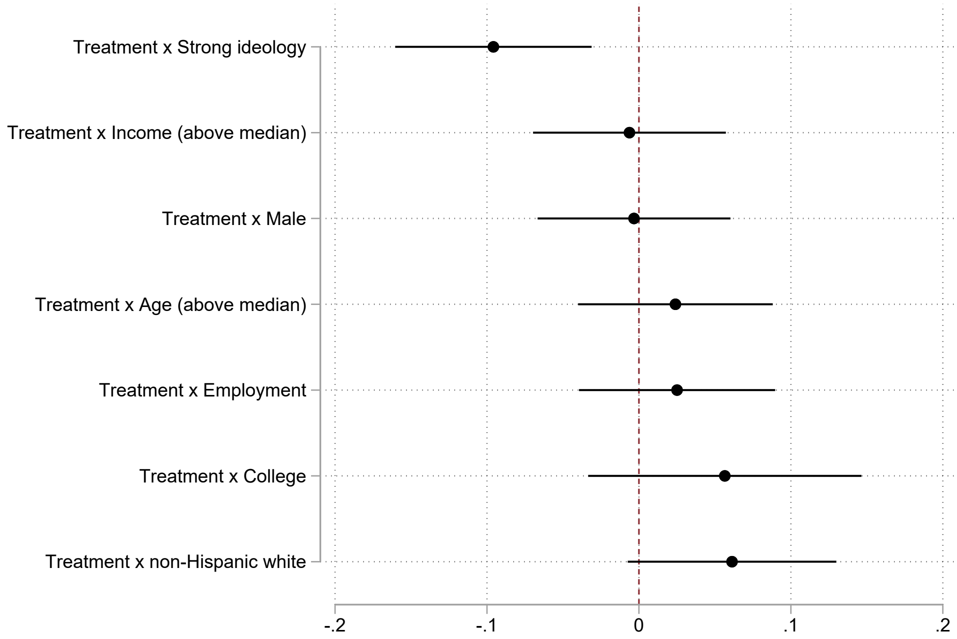

Figure 2.B.7. Heterogeneity in treatment effects on newsletter demand with MSNBC: Simultaneous interactions

Notes: This figure plots interaction coefficients (β_2) from a regression including our fact-check treatment, a vector of demographic controls and their interaction with the treatment indicator, i.e., a regression of the form $y = \beta_0 + \beta_1 \mathrm{Tr} + \beta_2 \mathrm{Tr} \times \boldsymbol{X}_i + \beta_3 \boldsymbol{X}_i + \varepsilon_i$ where $\boldsymbol{X}_i$ is a vector of demographic variables. 95% confidence intervals are indicated. The regression includes respondents who passed the attention check and were offered a newsletter featuring articles from *MSNBC*. "Strong ideology" is a binary variable taking value one for respondents who identify as "very liberal." "Income (above median)" is a binary variable taking value one if a respondent has above-median income. "Male" is a binary variable taking value one if a respondent is male. "Age (above median)" is a binary variable taking value one if a respondent has above-median age. "Employment" is a binary variable taking value one if a respondent is a full-time employee. "College" is a binary variable taking value one if a respondent has at least some college experience. "non-Hispanic White" is a binary variable taking value one if a respondent selected "Caucasian/White" and is of non-Hispanic origin.

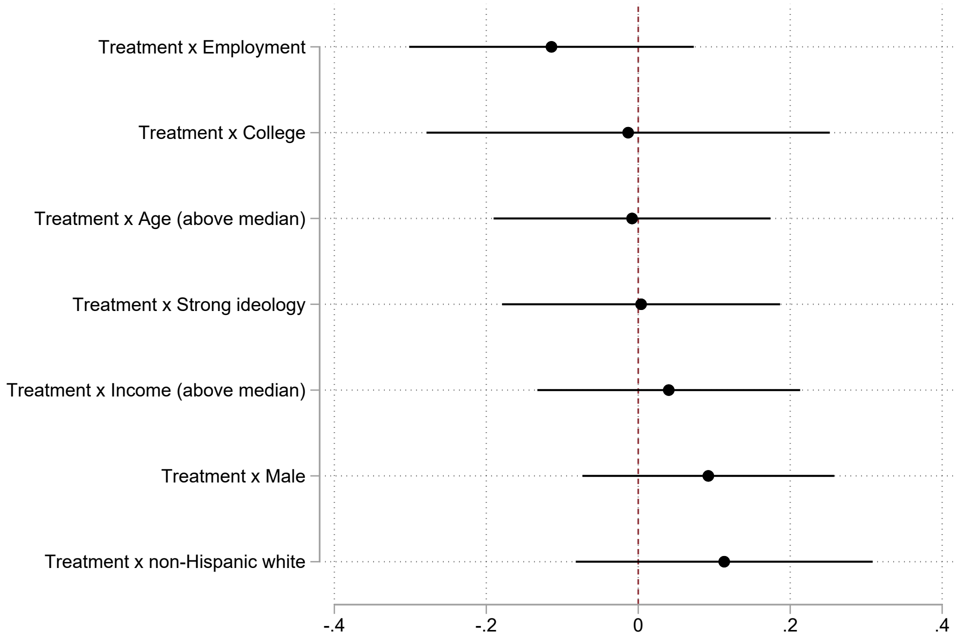

Figure 2.B.8. Heterogeneity in treatment effects on newsletter demand with Fox News: Simultaneous interactions

Notes: This figure plots interaction coefficients (β_2) from a regression including our fact-check treatment, a vector of demographic controls and their interaction with the treatment indicator, i.e., a regression of the form $y = \beta_0 + \beta_1 \mathrm{Tr} + \beta_2 \mathrm{Tr} \times \boldsymbol{X}_i + \beta_3 \boldsymbol{X}_i + \varepsilon_i$ where $\boldsymbol{X}_i$ is a vector of demographic variables. 95% confidence intervals are indicated. The regression includes respondents who passed the attention check and were offered a newsletter featuring articles from *Fox News*. “Strong ideology” is a binary variable taking value one for respondents who identify as “very liberal.” “Income (above median)” is a binary variable taking value one if a respondent has above-median income. “Male” is a binary variable taking value one if a respondent is male. “Age (above median)” is a binary variable taking value one if a respondent has above-median age. “Employment” is a binary variable taking value one if a respondent is a full-time employee. “College” is a binary variable taking value one if a respondent has at least some college experience. “non-Hispanic White” is a binary variable taking value one if a respondent selected “Caucasian/White” and is of non-Hispanic origin.

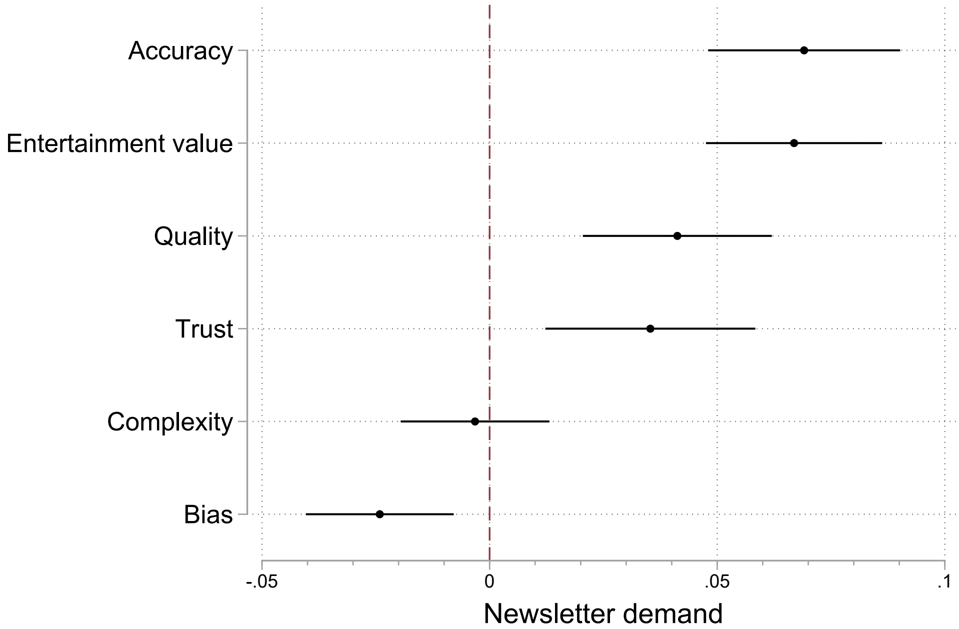

Figure 2.B.9. Correlates of demand: MSNBC

Notes: This figure plots the correlations between newsletter demand and a battery of z-scored beliefs about the newsletter from a joint regression that also controls for demographic characteristics. We use control group respondents that were offered a newsletter featuring articles from *MSNBC*. 95% confidence intervals are indicated.

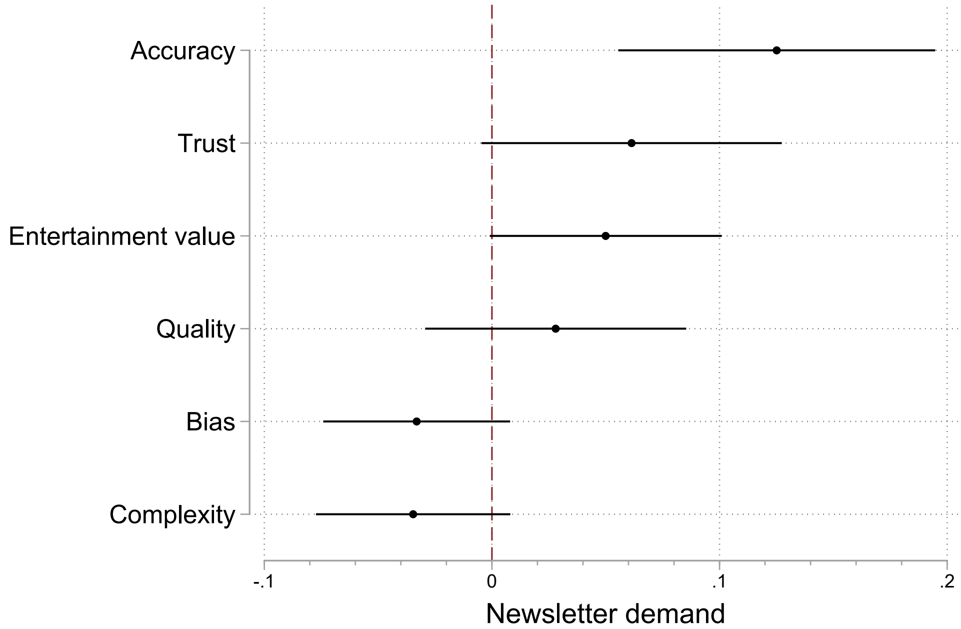

Figure 2.B.10. Correlates of demand: Fox News

Notes: This figure plots the correlations between newsletter demand and a battery of z-scored beliefs about the newsletter from a joint regression that also controls for demographic characteristics. We use control group respondents that were offered a newsletter featuring articles from *Fox News*. 95% confidence intervals are indicated.

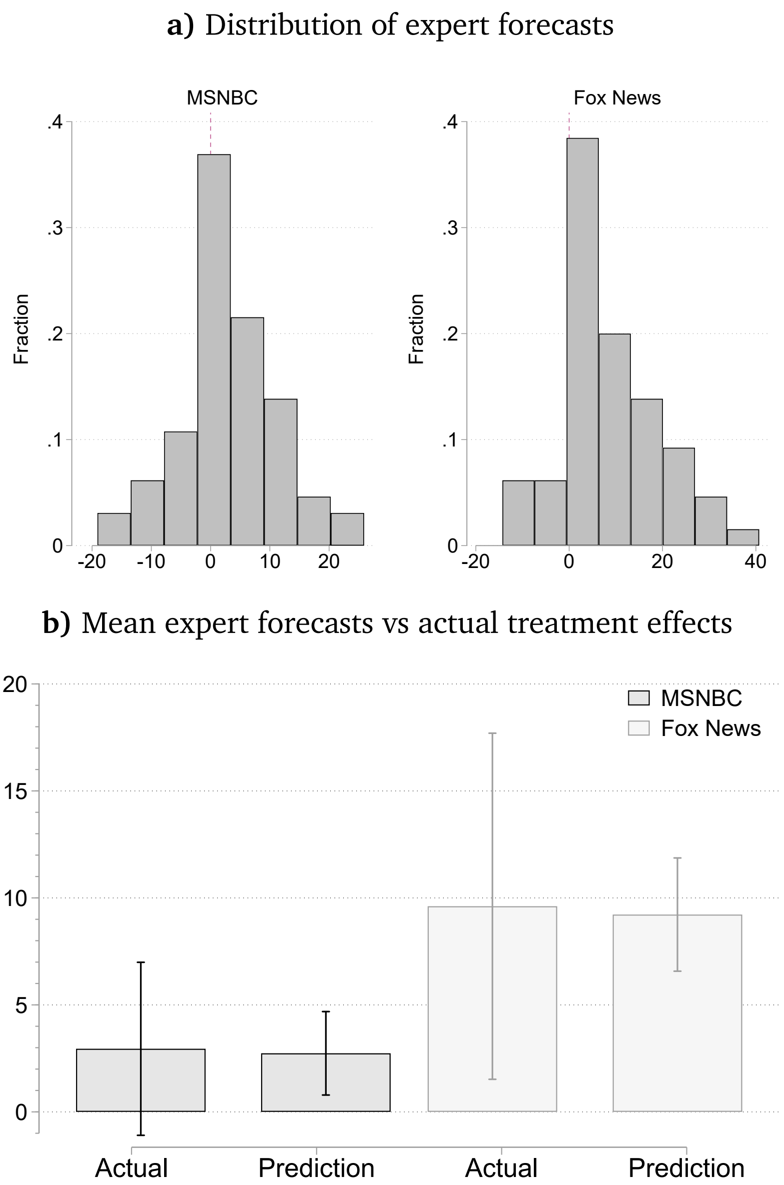

Figure 2.B.11. Expert survey

Notes: This figure uses data from the expert survey. Panel 2.B.11a shows the distribution of beliefs about treatment effects for *MSNBC* (left histogram) and *Fox News* (right histogram). Panel 2.B.11b shows the mean expert forecast of the treatment effects for *MSNBC* and *Fox News* contrasted with the actual treatment effects from the base treatments pooled across waves (estimated without controls but with wave fixed effects). 95% confidence intervals are indicated.

Appendix 2.C Comparing attentive and inattentive respondents

In this section, we compare respondents who passed our simple pre-treatment attention check (attentive respondents) and those who did not pass the attention check (inattentive respondents).[18] As shown below, there are several pieces of evidence indicating lower data quality among inattentive respondents:

- Given our sample of Biden voters, we would expect baseline demand for the newsletter featuring stories from the left-oriented *MSNBC* to be much higher than for the newsletter featuring stories from *Fox News*, a right-wing outlet. Among attentive respondents, baseline demand for the newsletter featuring stories from *MSNBC* is indeed 45% higher than for the newsletter featuring stories from *Fox News*. Among inattentive respondents, however, the difference in baseline demand is only 10.8% higher for the newsletter featuring stories from *MSNBC*.
- The median response time is 49 seconds higher for attentive respondents than for inattentive respondents. This corresponds to a 21.7% difference compared to the median response time of 226 seconds among inattentive respondents.[19] The significantly lower time spent on the survey is consistent with inattentive respondents not paying careful attention to details of the instructions.
- In Table 2.C.2, we display treatment effects of the fact-checking treatment on demand for the *MSNBC* newsletter and beliefs about newsletter characteristics separately for attentive and inattentive respondents. Panel B shows a large and significant first stage on beliefs about newsletter characteristics among attentive respondents. By contrast, Panel C shows that inattentive respondents do not adjust their beliefs about the characteristics of the newsletter.
- As shown in Figure 2.C.1, the correlations between newsletter demand and beliefs about newsletter characteristics—such as accuracy, quality, and trust—are much more pronounced in the control group sample of attentive respondents compared to inattentive respondents.

18. See page 180 for a screenshot of the attention check.
19. Table 2.C.1 shows similar patterns for average response time.

Table 2.C.1. Summary statistics: Full sample with attentive vs inattentive respondents

	(1) All respondents	(2) Attentive	(3) Inattentive
Male	0.440	0.400	0.488
Age	40.033	43.267	35.989
White	0.661	0.765	0.530
Log income	10.754	10.834	10.654
College education	0.810	0.864	0.742
Full-time employee	0.481	0.447	0.524
Northeast	0.235	0.229	0.242
Midwest	0.228	0.225	0.230
West	0.198	0.221	0.169
South	0.340	0.324	0.359
Hispanic	0.156	0.102	0.223
Time spent on survey	379.941	402.530	351.693
Demand: MSNBC	0.558	0.504	0.626
Demand: Fox News	0.481	0.389	0.589
Observations	8,399	4,667	3,732

Notes: This table displays the mean value of basic covariates for the full sample (column 1) and separately by whether respondents passed or did not pass a basic pre-treatment attention check (columns 2 and 3, respectively). "Male" is a binary variable with value one for male respondents. "Age" is age of the respondent. "White" is a binary variable with value one if the respondent selected "Caucasian/White." "Log income" is coded continuously as the log of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "College education" is a binary variable taking value one if the respondent selected "Some college, no degree," "Associates degree," "Bachelor's degree," or "Post-graduate degree." "Full-time employee" is a binary variable taking value one if the respondent is a full-time employee. "Northeast," "Midwest," "West" and "South" are binary variables with value one if the respondent lives in the respective region. "Hispanic" is a binary variable with value one if the respondent is Hispanic. "Time spent on survey" is the number of seconds the respondents spent on the survey. "Demand: MSNBC" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter featuring stories from *MSNBC* and zero for respondents who said "No." "Demand: Fox News" is similarly defined for respondents featured the newsletter featuring stories from *Fox News*.

Table 2.C.2. Heterogeneity by attention: MSNBC

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Full sample							
Treatment	0.003	0.068***	0.018	0.021	-0.020	0.014	0.001
	(0.011)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.022)
N	7,371	7,294	7,294	7,294	7,294	7,294	7,294
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.556	0	0	0	0	0	0
Panel B: Attentive							
Treatment	0.014	0.143***	0.087***	0.049	-0.051*	0.035	0.023
	(0.016)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.030)
N	4,109	4,069	4,069	4,069	4,069	4,069	4,069
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.497	0	0	0	0	0	0
Panel C: Inattentive							
Treatment	-0.009	-0.002	-0.051	-0.000	0.007	-0.005	-0.023
	(0.016)	(0.035)	(0.036)	(0.034)	(0.034)	(0.035)	(0.034)
N	3,262	3,225	3,225	3,225	3,225	3,225	3,225
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.631	0	0	0	0	0	0

Notes: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use respondents that were offered a newsletter featuring *MSNBC* articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents who passed the attention check. Panel C shows results respondents who did not pass the attention check. "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased.." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining."

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2.C.3. Heterogeneity by attention: Fox News

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Full sample							
Treatment	0.079***	0.163***	0.165***	0.177***	-0.140**	-0.113*	0.144**
	(0.030)	(0.060)	(0.060)	(0.061)	(0.060)	(0.062)	(0.061)
N	1,028	1,010	1,010	1,010	1,010	1,010	1,010
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.445	0	0	0	0	0	0
Panel B: Attentive							
Treatment	0.100**	0.231***	0.152*	0.177**	-0.124	-0.076	0.107
	(0.041)	(0.084)	(0.084)	(0.085)	(0.081)	(0.086)	(0.087)
N	558	548	548	548	548	548	548
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.343	0	0	0	0	0	0
Panel C: Inattentive							
Treatment	0.048	0.057	0.156*	0.151*	-0.170*	-0.146	0.168*
	(0.044)	(0.088)	(0.086)	(0.088)	(0.088)	(0.091)	(0.086)
N	470	462	462	462	462	462	462
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.569	0	0	0	0	0	0

Notes: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use respondents that were offered a newsletter featuring *Fox News* articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents who passed the attention check. Panel C shows results respondents who did not pass the attention check. "Treatment" is a binary variable taking value one if the articles in the newsletter are fact-checked. "News demand" is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No." "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." "Left-wing bias" is measured on a 5-point scale from "Very right-wing biased" to "Very left-wing biased.." "Complexity" of the newsletter articles is measured on a 5-point scale from "Very simple" to "Very complex." "Entertainment" of the newsletter is measured on a 5-point scale from "Not entertaining at all" to "Very entertaining."

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

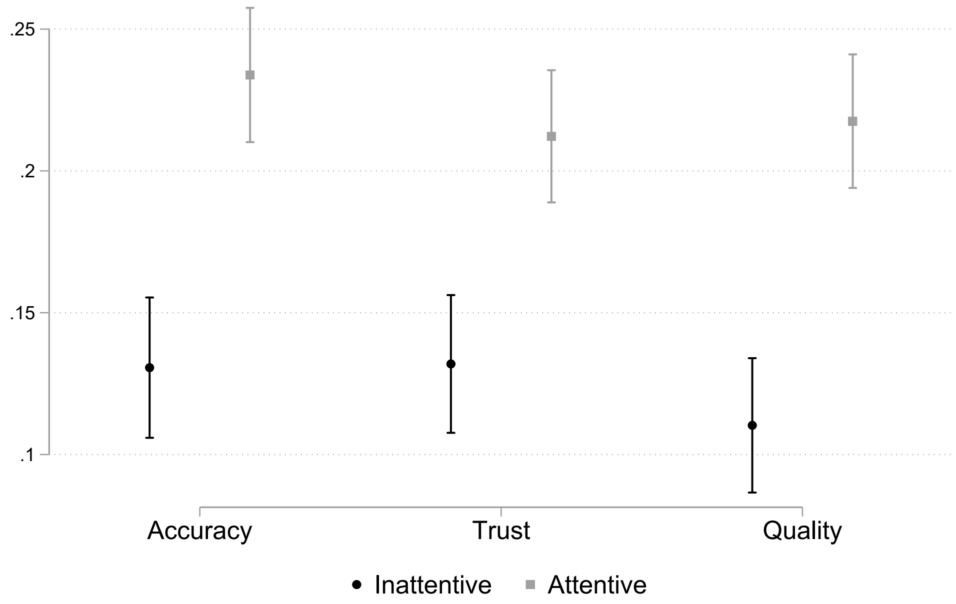

Figure 2.C.1. Correlates of newsletter demand by attention

Notes: This figure shows coefficient plots from bivariate OLS regressions where the dependent variable is newspaper demand (which is a binary variable taking the value one for respondents who said "Yes" to receive the newsletter and zero for those who said "No"). The independent variables are different beliefs about the newsletter characteristics (accuracy, trust, and quality). "Accuracy" of the newsletter is measured on a 5-point scale from "Very inaccurate" to "Very accurate." "Trust" is the trustworthiness of the newsletter and measured on a 5-point scale from "Not trustworthy at all" to "Very trustworthy." "Quality" of the newsletter is measured on a 5-point scale from "Very low quality" to "Very high quality." All regressions use control group respondents that were offered a newsletter featuring articles from *MSNBC*. We run the regressions separately for respondents who passed and did not pass the pre-treatment attention check. 95% confidence intervals are indicated.

Appendix 2.D Fact-checking

While we did not explicitly reveal to our respondents how we selected the three top stories, in practice we used Google News to identify the three top stories about the *Biden Rescue Plan* from *MSNBC* and *Fox News*. We then employed two complementary approaches to fact-check the veracity of the information contained in featured articles. First, we fact-checked the articles using the following steps:

- Identify whether a similar news article appeared in other high-quality outlets (e.g. Reuters). Then search for inconsistencies across these articles.
- Identify the primary source of statistical information, assess whether they are accurately represented, and compare the figures to estimates from other, high-quality sources (e.g. government reports, published studies).
- Identify the primary source of quotations and assess whether they are quoted out of context.

Second, we collected information on inaccurate claims from well-known fact-checking organizations to rule out that we missed already identified false claims. Below we provide two examples of false claims.

MSNBC. On March 12, 2021, *MSNBC* published the article "Dems' COVID relief package already saving tens of thousands of jobs." In this article, the author claims that independent economic forecasts have "projected the law may create as many as 7 million jobs," citing a projection by Gregory Daco. This is misleading because the projection includes both the effect of the fiscal stimulus as well as improving economic conditions. This example illustrates how the ideologically aligned outlet biased their reports towards the beliefs of their readers by making exaggerated claims about the positive consequences of the stimulus plan.

Fox News. On March 7, 2021, *Fox News* published the article "Sen. Blackburn on massive coronavirus package heading to House without GOP support." This article focuses on the critique of Senator Marsha Blackburn that "only nine percent" of the spending involved in the stimulus plan is related to fighting the coronavirus. While spending on vaccines and other medical supplies accounts for about nine percent, the stimulus plan also includes financial relief for households affected by the pandemic.

Table 2.D.1 below provides a screenshot of the website where we published our newsletter. The release schedule for our politics newsletters is shown in Table 2.D.1. As the *Biden Rescue Plan* was signed into law on March 11, 2021, we ceased to publish weekly updated at this point.

Newsletter

Biden Rescue Plan

Thank you very much for your interest in our weekly newsletter.

What is the newsletter about?

The newsletter will cover the **Biden Rescue Plan**, which is a plan to help America recover from the coronavirus crisis. The newsletter will contain three top stories about the plan featured on **MSNBC** during the last week.

How do I get the newsletter?

We will publish our weekly politics newsletter each Monday in January and Feburary 2021 on this website.

Latest issues

Top articles from **January**:

- What Biden's economic relief plan gets right
- Can Joe Biden's economic rescue plan get through Congress?
- Why Republicans are 'frustrated' and 'angry' as Dems eye COVID relief

Top articles from **Feburary**

- Republicans' Covid-19 stimulus compromise gives Biden little to work with
- GOP lawmakers push back against expanded unemployment assistance
- Slowly but surely, COVID relief bill inches forward, gains support

Figure 2.D.1. Newsletter about the *Biden Rescue Plan*

Notes: This is a screenshot of the website where we published our newsletter.

Table 2.D.1. Release schedule of the politics newsletter

			Wave 1	**Wave 2**	**Wave 3**	**Wave 4**
Month	Day of month	Week number	Topic polarization	Instrumental motives	Fox News	Opinion
January	4	1				
	11	2				
	18	3				
	25	4	X	X		
February	1	5	X	X		
	8	6	X	X		
	15	7	X	X		
	22	8	X	X	X	X
March	1	9	X	X	X	X
	8	10	X	X	X	X
	15	11	Biden Rescue Plan is signed into law at this point			

Notes: This table shows the release schedule of our newsletter for each wave. Both wave 1 and wave 2 used the same set of articles. The Biden Rescue Plan was signed into law on March 11, 2021. On March 15, the newsletter informed recipients about this fact and announced that it would cease to publish weekly updates. At the end of March, we deactivated the newsletter websites.

Appendix 2.E Screenshots

2.E.1 Full survey with base treatments (identical across all waves)

2.E.1.1 Pre-treatment questions

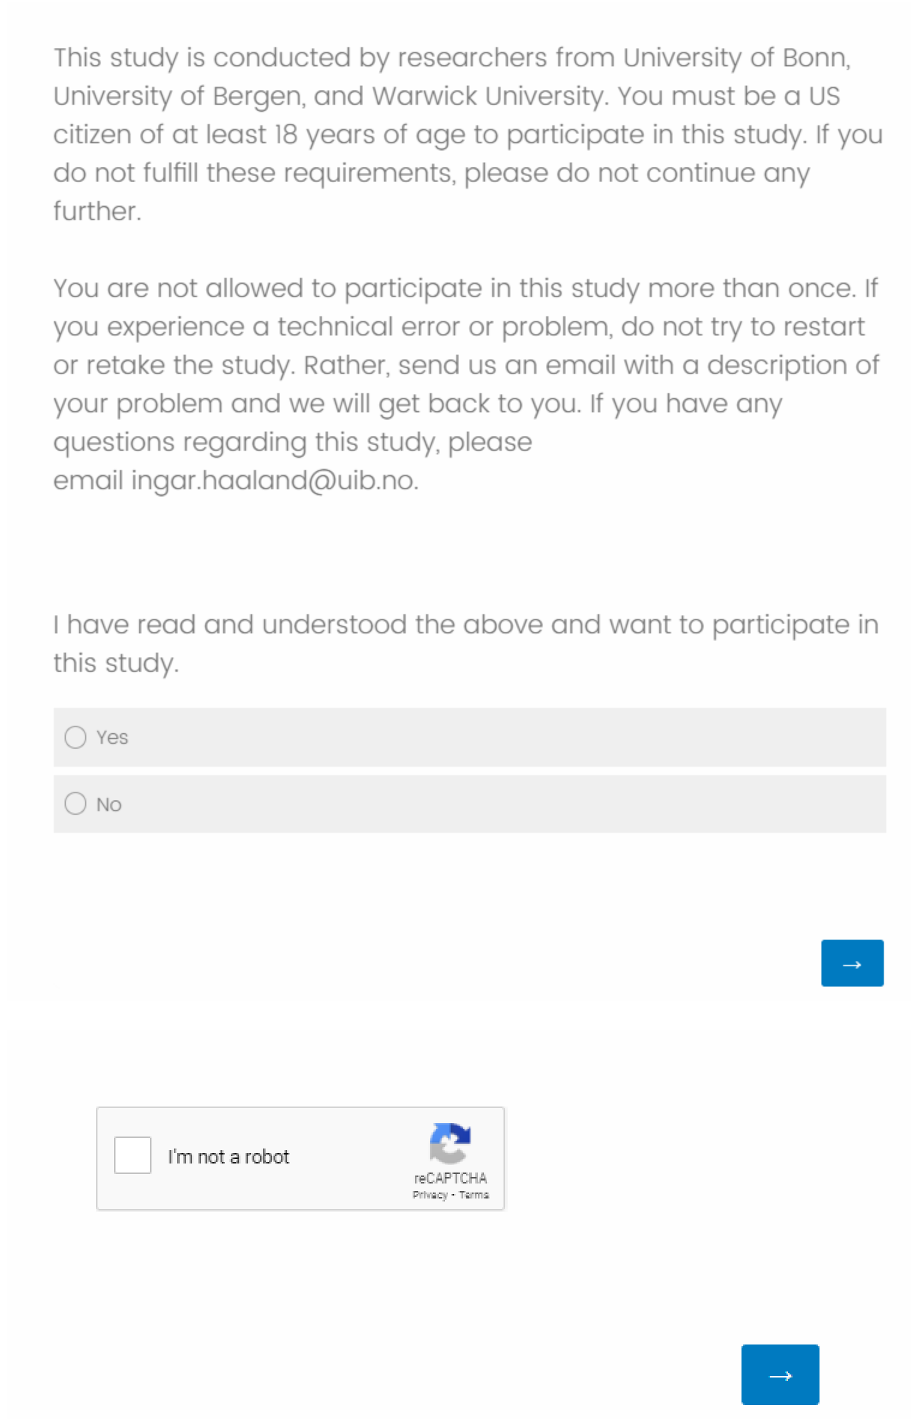

This study is conducted by researchers from University of Bonn, University of Bergen, and Warwick University. You must be a US citizen of at least 18 years of age to participate in this study. If you do not fulfill these requirements, please do not continue any further.

You are not allowed to participate in this study more than once. If you experience a technical error or problem, do not try to restart or retake the study. Rather, send us an email with a description of your problem and we will get back to you. If you have any questions regarding this study, please email ingar.haaland@uib.no.

I have read and understood the above and want to participate in this study.

- Yes
- No

→

I'm not a robot

reCAPTCHA

Privacy - Terms

→

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This means that there are a lot of random answers which compromise the results of research studies. To show that you read our questions carefully, please choose both "Extremely interested" and "Not at all interested" as your answer in the next question. How interested are you in sports?

- ☐ Extremely interested
- ☐ Very interested
- ☐ A little bit interested
- ☐ Almost not interested
- ☐ Not at all interested

→

Please indicate your gender.

- ◯ Male
- ◯ Female

What is your age?

☐

Which category best describes your highest level of education?

- ◯ Eighth grade or less
- ◯ Some high school
- ◯ High school degree/GED
- ◯ Some college
- ◯ 2-year college degree
- ◯ 4-year college degree
- ◯ Master's degree
- ◯ Doctoral degree
- ◯ Professional degree (JD, MD, MBA)

→

Which of the following best describes your race or ethnicity?

- ◯ African American/Black
- ◯ Asian/Asian American
- ◯ Caucasian/White
- ◯ Native American, Inuit or Aleut
- ◯ Native Hawaiian/Pacific Islander
- ◯ Other

Are you of Hispanic, Latino, or Spanish origin?

- ◯ Yes
- ◯ No

What was your family's gross household income in 2020 in US dollars?

- ◯ Less than $15,000
- ◯ $15,000 to $24,999
- ◯ $25,000 to $49,999
- ◯ $50,000 to $74,999
- ◯ $75,000 to $99,999
- ◯ $100,000 to $149,999
- ◯ $150,000 to $200,000
- ◯ More than $200,000

Who did you vote for in the 2020 presidential election?

- ◯ Donald Trump
- ◯ Joe Biden
- ◯ Other
- ◯ Did not vote

In politics, as of today, do you consider yourself a Republican, a Democrat, or an Independent?

- ◯ Republican
- ◯ Democrat
- ◯ Independent

→

What is your region of residence?

- ◯ **Northeast** (CT, ME, MA, NH, RI, VT, NJ, NY,PA),
- ◯ **Midwest** (IL, IN, MI, OH, WI, IA, KS, MN, MO, NE, ND, SD)
- ◯ **South** (DE, DC, FL, GA,MD, NC, SC, VA, WV, AL, KY, MS, TN, AR, LA, OK, TX)
- ◯ **West** (AZ, CO, ID, NM, MT, UT,NV, WY, AK, CA, HI, OR, WA)

What is your current employment status?

- ◯ Full-time employee
- ◯ Part-time employee
- ◯ Self-employed or small business owner
- ◯ Unemployed and looking for work
- ◯ Student
- ◯ Not in labor force (for example: retired or full-time parent)

Are you liberal or conservative?

- ◯ Very liberal
- ◯ Liberal
- ◯ Neither liberal nor conservative
- ◯ Conservative
- ◯ Very conservative

→

Which of the following newspapers are you most likely to read?

- Breitbart
- BuzzFeed News
- Chicago Sun-Times
- Daily Mail
- Drudge Report
- InfoWars
- Los Angeles Times
- New Republic
- Newsmax
- New York Daily News
- New York Post
- Palmer Report
- The Denver Post
- The Huffington Post
- The Mercury News
- The New York Times
- The Wall Street Journal
- The Washington Post
- The Washington Times
- USA Today
- I never read any of the newspapers above

→

2.E.1.2 Newsletter without fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in January and February 2021.

Would you like to receive our newsletter?

- ◯ Yes
- ◯ No

→

2.E.1.3 Newsletter with fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in January and February 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

◯ Yes

◯ No

→

2.E.1.4 Post-treatment mechanism questions

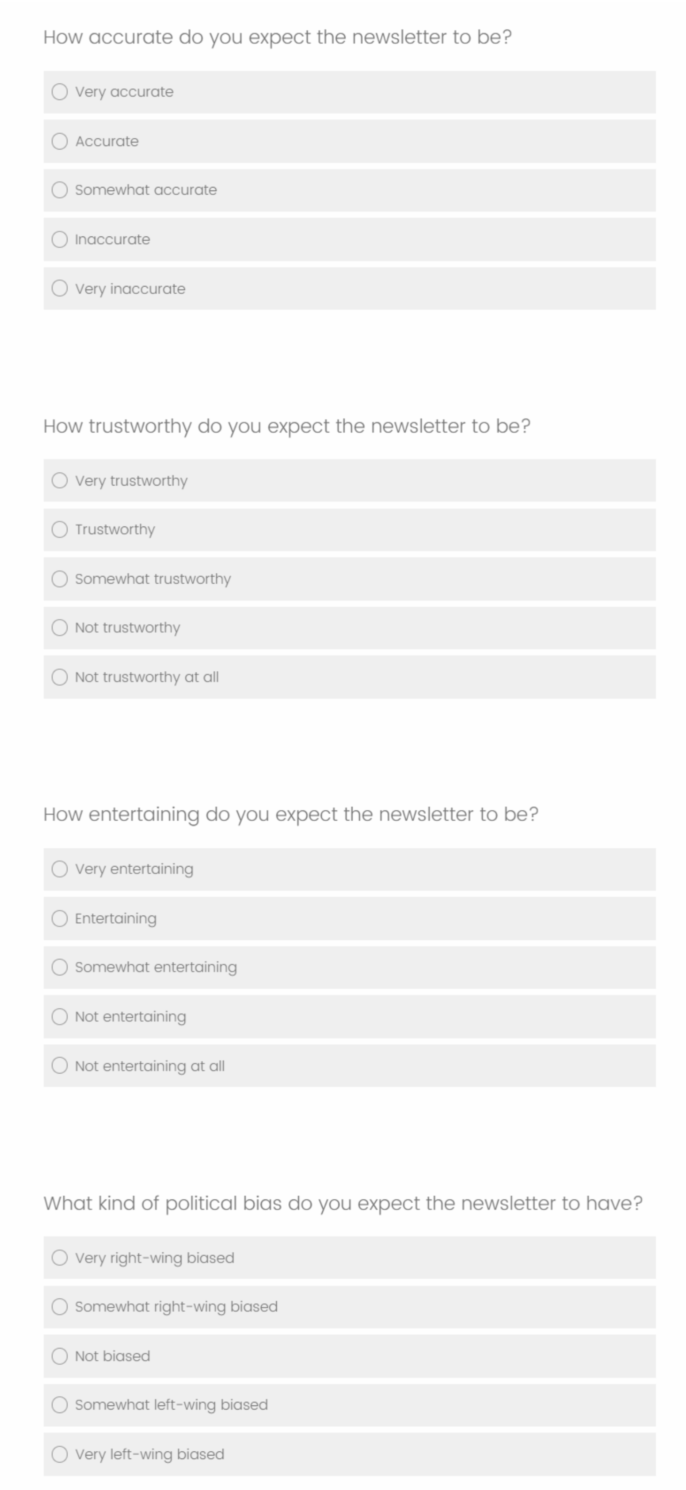

How accurate do you expect the newsletter to be?

- Very accurate
- Accurate
- Somewhat accurate
- Inaccurate
- Very inaccurate

How trustworthy do you expect the newsletter to be?

- Very trustworthy
- Trustworthy
- Somewhat trustworthy
- Not trustworthy
- Not trustworthy at all

How entertaining do you expect the newsletter to be?

- Very entertaining
- Entertaining
- Somewhat entertaining
- Not entertaining
- Not entertaining at all

What kind of political bias do you expect the newsletter to have?

- Very right-wing biased
- Somewhat right-wing biased
- Not biased
- Somewhat left-wing biased
- Very left-wing biased

What quality would you expect the newsletter to have?

- ◯ Very high quality
- ◯ High quality
- ◯ Medium quality
- ◯ Low quality
- ◯ Very low quality

Do you expect the newsletter to have a simple or complex message?

- ◯ Very simple
- ◯ Simple
- ◯ Neither simple nor complex
- ◯ Complex
- ◯ Very complex

How much trust do you have in the news media?

- ◯ Very high trust
- ◯ High trust
- ◯ Some trust
- ◯ Low trust
- ◯ Very low trust

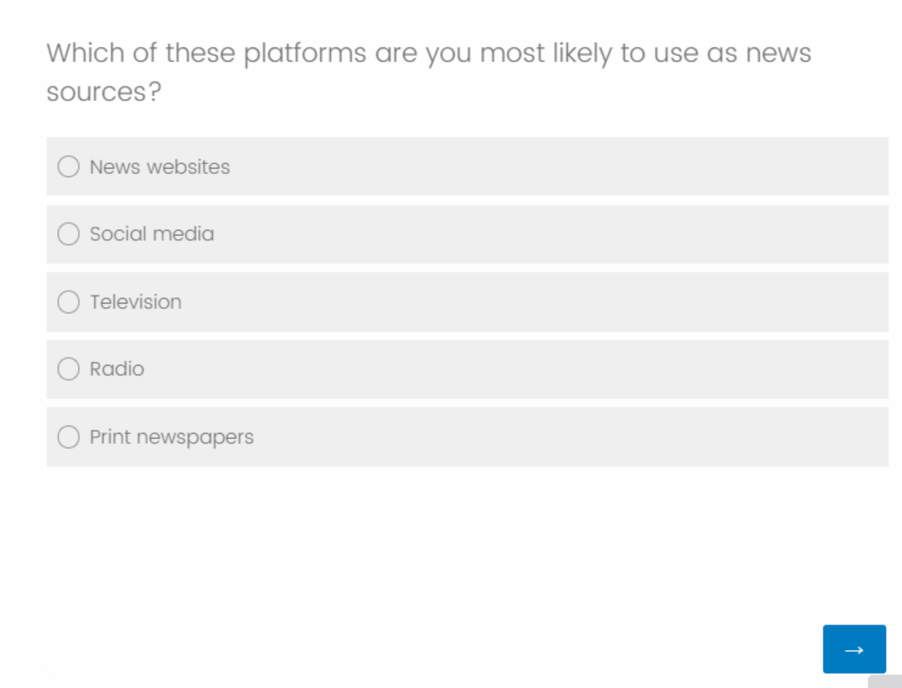

Which of these platforms are you most likely to use as news sources?

- ◯ News websites
- ◯ Social media
- ◯ Television
- ◯ Radio
- ◯ Print newspapers

→

2.E.1.5 Beliefs about fact-checking: condition 1

Now imagine that we would fact check all stories and flag those with inaccuracies.

How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?

- ○ 0
- ○ 1
- ○ 2
- ○ 3

How many of the top three articles from MSNBC selected for the newsletter do you expect to be flagged as inaccurate?

- ○ 0
- ○ 1
- ○ 2
- ○ 3

How much do you trust our ability to fact check articles from MSNBC?

- ○ Strongly trust
- ○ Trust
- ○ Somewhat trust
- ○ Do not trust
- ○ Do not trust at all

→

2.E.1.6 Beliefs about fact-checking: condition 2

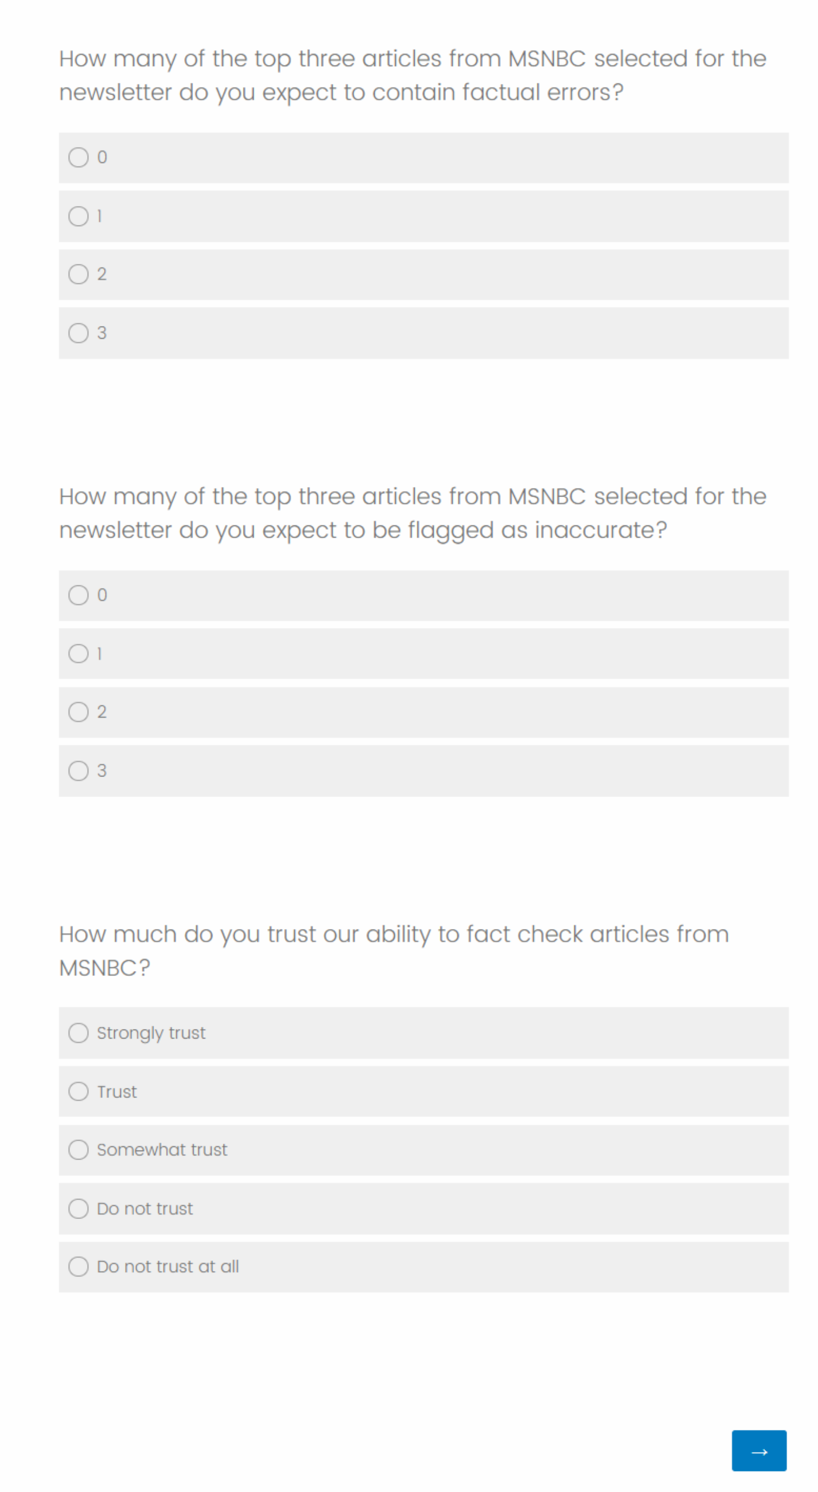

How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?

- ○ 0
- ○ 1
- ○ 2
- ○ 3

How many of the top three articles from MSNBC selected for the newsletter do you expect to be flagged as inaccurate?

- ○ 0
- ○ 1
- ○ 2
- ○ 3

How much do you trust our ability to fact check articles from MSNBC?

- ○ Strongly trust
- ○ Trust
- ○ Somewhat trust
- ○ Do not trust
- ○ Do not trust at all

→

2.E.1.7 Demand for fact-checking information

MSNBC has been fact-checked several times by non-partisan fact checkers. Do you want to know how the accuracy of MSNBC has been rated?

- ◯ Yes
- ◯ No

→

2.E.1.8 Questions about the Biden Rescue Plan

How strongly do you support the Biden Rescue Plan?

- ◯ Strongly support
- ◯ Support
- ◯ Neither support nor oppose
- ◯ Oppose
- ◯ Strongly oppose

→

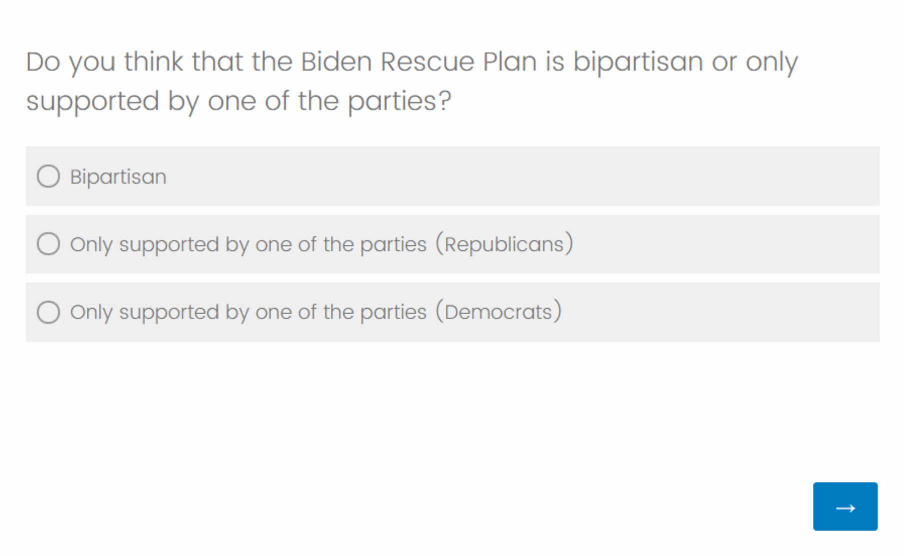

Do you think that the Biden Rescue Plan is bipartisan or only supported by one of the parties?

- Bipartisan
- Only supported by one of the parties (Republicans)
- Only supported by one of the parties (Democrats)

→

2.E.2 Wave 1: Topic polarization

2.E.2.1 Newsletter without fact-checking

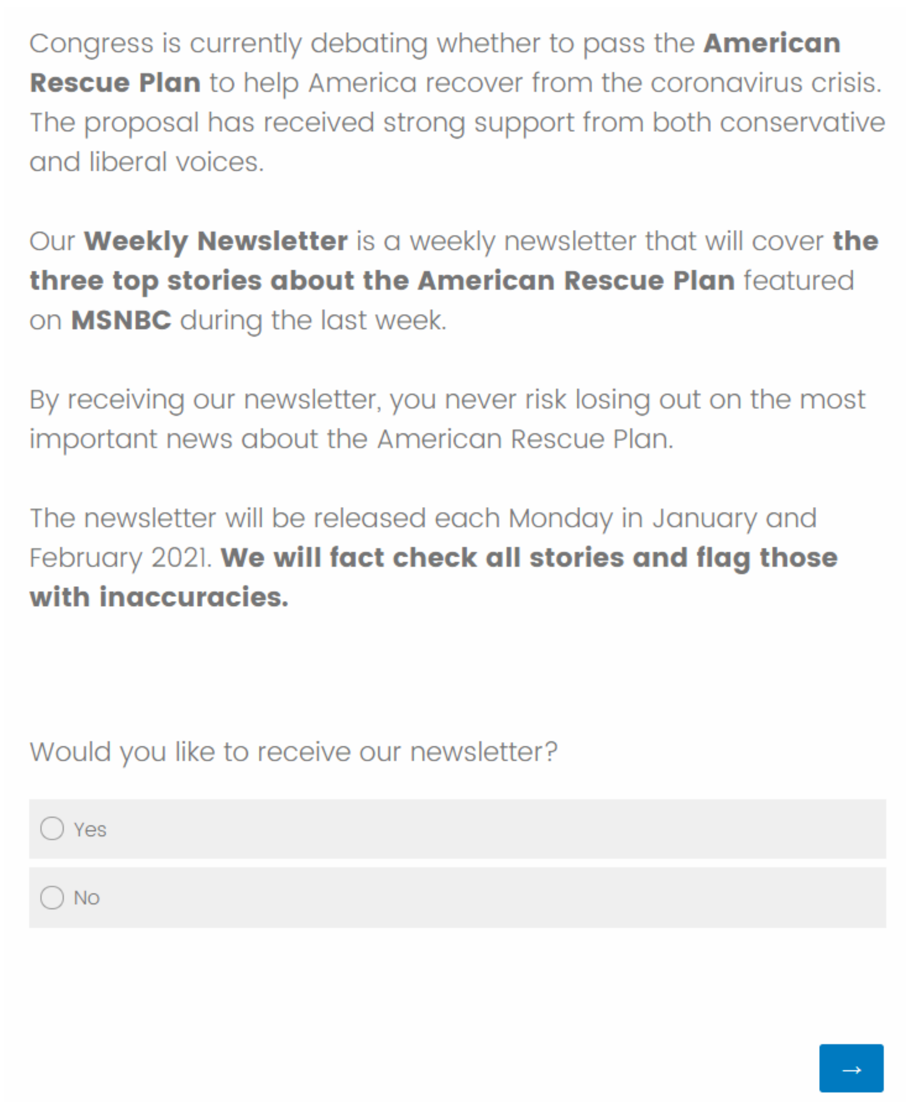

Congress is currently debating whether to pass the **American Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from both conservative and liberal voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the American Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the American Rescue Plan.

The newsletter will be released each Monday in January and February 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

- Yes
- No

→

2.E.2.2 Newsletter with fact-checking

Congress is currently debating whether to pass the **American Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from both conservative and liberal voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the American Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the American Rescue Plan.

The newsletter will be released each Monday in January and February 2021.

Would you like to receive our newsletter?

- ○ Yes
- ○ No

→

2.E.3 Wave 2: Instrumental motives

2.E.3.1 Newsletter without fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan,** which includes a **$1400 stimulus check to most Americans**, to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan, including the **latest news about when you could you get your $1,400 if the plan is approved.**

The newsletter will be released each Monday in January and February 2021.

Would you like to receive our newsletter?

◯ Yes

◯ No

→

2.E.3.2 Newsletter with fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan,** which includes a **$1400 stimulus check to most Americans**, to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan, including the **latest news about when you could you get your $1,400 if the plan is approved.**

The newsletter will be released each Monday in January and February 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

- Yes
- No

→

2.E.3.3 Manipulation checks for instrumental motives

How relevant do you expect the newsletter to be for your personal finances?

- ◯ Very relevant
- ◯ Relevant
- ◯ Somewhat relevant
- ◯ Not relevant
- ◯ Not relevant at all

How do you expect the Biden Rescue Plan to affect your personal finances?

- ◯ Affect my personal finances very positively
- ◯ Affect my personal finances positively
- ◯ Does not affect my personal finances
- ◯ Affect my personal finances negatively
- ◯ Affect my personal finances very negatively

2.E.4 Wave 3: Right-wing outlet

2.E.4.1 Newsletter without fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **Fox News** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in February and March 2021.

Would you like to receive our newsletter?

○ Yes

○ No

→

2.E.4.2 Newsletter with fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top stories about the Biden Rescue Plan** featured on **Fox News** during the last week.

By receiving our newsletter, you never risk losing out on the most important news about the Biden Rescue Plan.

The newsletter will be released each Monday in February and March 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

- ○ Yes
- ○ No

→

2.E.5 Wave 4: Opinion piece

2.E.5.1 Newsletter without fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top opinion pieces about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most newsworthy opinion pieces about the Biden Rescue Plan.

The newsletter will be released each Monday in February and March 2021.

Would you like to receive our newsletter?

○ Yes

○ No

→

2.E.5.2 Newsletter with fact-checking

Congress is currently debating whether to pass the **Biden Rescue Plan** to help America recover from the coronavirus crisis. The proposal has received strong support from liberal voices, but has been criticized by conservative voices.

Our **Weekly Newsletter** is a weekly newsletter that will cover **the three top opinion pieces about the Biden Rescue Plan** featured on **MSNBC** during the last week.

By receiving our newsletter, you never risk losing out on the most newsworthy opinion pieces about the Biden Rescue Plan.

The newsletter will be released each Monday in February and March 2021. **We will fact check all stories and flag those with inaccuracies.**

Would you like to receive our newsletter?

- ◯ Yes
- ◯ No

→

Chapter 3

Fighting Climate Change: The Role of Norms, Preferences, and Moral Values

Joint with Peter Andre, Teodora Boneva, and Armin Falk

Abstract: We document individual willingness to fight climate change and its behavioral determinants in a large representative sample of US adults. Willingness to fight climate change – as measured through an incentivized donation decision – is highly heterogeneous across the population. Individual beliefs about social norms, economic preferences such as patience and altruism, as well as universal moral values positively predict climate preferences. Moreover, we document systematic misperceptions of prevalent social norms. Respondents vastly underestimate the prevalence of climate-friendly behaviors and norms among their fellow citizens. Providing respondents with correct information causally raises individual willingness to fight climate change as well as individual support for climate policies. The effects are strongest for individuals who are skeptical about the existence and threat of global warming.

Acknowledgements: We thank Chris Roth for helpful comments and discussions; Ana Bras Monteiro and Tobias Reinheimer for excellent research assistance; and Markus Antony for administrative support. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866. Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A01, B03) is gratefully acknowledged. The study obtained ethics approval from the German Association for Experimental Economic Research (#Xx5i4FQa, 02/09/2021). The main research questions, the survey design, and the sampling approach were preregistered at the AEA RCT Registry (#AEARCTR-0007542). Data and code will be made available.

3.1 Introduction

Climate change is one of the greatest threats facing humanity today. Its social and economic implications range from increased mortality and violence to reduced human productivity and economic growth (IPCC, 2014; Carleton and Hsiang, 2016; Auffhammer, 2018). The estimated economic impacts are enormous. Studies indicate that climate change could lower global GDP by 23% by 2100 and further exacerbate existing inequalities (Burke, Hsiang, and Miguel, 2015; Diffenbaugh and Burke, 2019). While many countries around the world have committed to meeting the 1.5 or 2 degree targets set out in the Paris Agreement, progress towards these goals has been slow (UNEP, 2019). In fact, it has become increasingly likely that global temperatures may rise well above the 2 degree target throughout the course of this century, with potentially catastrophic impacts for both human society and ecosystems. Given the threat posed by climate change, it is important to understand what determines people's individual willingness to fight climate change, namely their willingness to engage in climate-friendly, sustainable, but potentially costly behavior. Understanding the determinants of these individual 'climate preferences' can help us to design effective policies against climate change that complement existing regulatory frameworks, such as carbon taxation.

In this paper, we shed light on the behavioral determinants of climate preferences. We explore the role of economic preferences, universal moral values, and beliefs about social norms. We also design a norm intervention to examine whether informing individuals about the prevalence of climate norms raises their willingness to fight climate change and their support for climate policies.

For this purpose, we administer a survey to a large representative sample of 8,000 US adults. We elicit individual willingness to fight climate change using an incentivized donation decision. More specifically, respondents are asked to divide $450 between themselves and a charitable organization that fights global warming. This decision captures the central trade-off that individuals face when deciding whether to take climate action, namely the notion that protecting the climate comes at a cost. To incentivize the decision, we implement the choices of a random subset of participants. The more money the respondents are willing to forgo and donate, the higher their willingness to fight climate change. To shed light on the potential determinants of climate preferences, we obtain detailed, individual-level information on perceived social norms, fundamental economic preferences, and moral values. We measure perceived social norms by asking respondents to estimate (i) the share of the US population that tries to fight global warming ('perceived behavior') and (ii) the share of the US population that thinks people in the US *should* try to fight global warming ('perceived norms'). To elicit economic preferences, we administer an experimentally validated survey to measure patience, willingness to take risks, altruism, trust, positive reciprocity, and negative reciprocity (Falk, Becker, Dohmen, Enke, Huffman, et al. (2018) and Falk, Becker, Dohmen, Huffman, and

Sunde (2018)). We further administer the Moral Foundations Questionnaire to obtain a measure of the relative importance of universal versus communal moral values (Haidt and Joseph (2004), Haidt (2012), Graham, Haidt, Koleva, Motyl, Iyer, et al. (2013), and Enke (2020)).

A natural question that arises is whether it is possible to raise individual willingness to fight climate change. While it is difficult to alter some behavioral determinants such as fundamental economic preferences or moral values, at least in the short run, beliefs about social norms are likely to be considerably more malleable. We therefore conduct a survey experiment to study the extent to which information provision can raise individual willingness to fight climate change. Respondents are randomized into a control condition or one of two treatments. The 'behavior treatment' provides respondents with truthful information about the proportion of the US population who try to fight global warming (62%), while the 'norms treatment' informs respondents about the true share of the US population who think that people in the US should try to fight global warming (79%). These low-cost information treatments have the potential to correct misperceptions about prevalent behaviors and norms and may shift individual willingness to fight global warming.

Several findings emerge from our study. First, we document large heterogeneity in individual willingness to fight climate change. In particular, climate preferences are systematically related to perceived social norms, economic preferences, as well as universal moral values. Conditional on a large set of covariates, perceived social norms strongly predict individual willingness to fight global warming. A one-standard-deviation increase in the perceived share of Americans trying to fight global warming is associated with a \$12 higher donation amount, while a corresponding increase in the perceived share of Americans who think that people in the US *should* try to fight global warming is associated with a \$14 higher donation. These results are consistent with individuals being 'conditional cooperators'. Put differently, respondents may be more willing to fight climate change if they believe that a higher proportion of their fellow citizens do the same. Among the economic preferences that we measure, patience, altruism, and positive reciprocity positively predict individual willingness to fight global warming. Similarly, universal moral values are positively associated with larger donations. Individuals with universal moral values are more willing to fight climate change compared to individuals who endorse communal, in-group-oriented values. The fight against climate change can be viewed as a global cooperation problem affecting present and future generations all around the world. It therefore is plausible that more patient and prosocial individuals as well as individuals with universal moral values more strongly value climate protection. Our finding that fundamental human traits, such as altruism, positive reciprocity, and moral universalism, are strong predictors of individual willingness to fight climate change helps us to understand the frequently observed cultural and political dissent on climate change (Dunlap, McCright, and Yarosh, 2016; Hornsey, Harris, and Fielding, 2018). In our data, economic preferences and universalism together explain about 40% of the large partisan gap in willingness to fight climate change.

Second, we document large heterogeneity in beliefs about prevalent behaviors and norms in the US. We find that respondents on average misperceive prevalent social norms. On average, respondents in our sample underestimate the true share of Americans who try to fight global warming as well as the true share of Americans who think that people in the US should try to fight global warming. This underestimation of climate norms is concerning because it could hamper individual willingness to fight climate change. Whether or not correcting these misperceptions can shift climate behavior is a question that we explore with the survey experiment.

Third, we find that both treatments positively affect individual willingness to fight climate change. Being informed about the true share of Americans who try to fight global warming raises donations by $12 (or 4.7%), while being informed about the true share of Americans who think that people in the US should try to fight global warming increases donations by $16 (or 6.3%). The effect sizes are strong considering the minimalist nature of the interventions. A heterogeneity analysis reveals that the positive treatment effects on the donation amount are primarily driven by the subgroup of respondents whose prior beliefs lie below the actual shares. Reassuringly, we do not observe a back-firing effect among respondents with prior beliefs above the actual shares. For them, the estimated treatment effects are also positive, albeit insignificant. We further explore whether the information treatments differentially affect individuals who are more or less skeptical about the existence and threat of human-caused climate change. We find that the information treatments are more effective for 'climate change deniers', who may have been surprised to learn that they hold minority views. The results are promising as they suggest that simple, low-cost informational interventions may be well-suited to reach skeptical subgroups of the population who are otherwise difficult to reach and convince.

Finally, we study whether the treatments causally affect individual support for climate policies (e.g. a carbon tax, subsidies for green energy, pollution regulation) and individual willingness to engage in political actions (e.g. volunteer time, attend a protest, contact government officials). Both treatments significantly raise individual support for climate policies. Again, the estimated treatment effects are stronger for the subgroup of the population who we classify as 'climate change deniers'.

Our findings have important implications for climate politics. Misperceptions of climate norms prevail in the US and can form a dangerous obstacle to climate action. However, at the same time, they can provide a unique opportunity to promote and accelerate climate-friendly behavior. A simple, easily scalable, and cost-effective intervention can correct these misperceptions and encourage climate-friendly behavior. This intervention is particularly effective for climate change skeptics, who are commonly difficult to reach but crucial for building up a broad alliance against climate change. Our results suggest that social norms should play a pivotal role in the policy response to climate change. Policies that foster social norms should complement formal regulations. For example, while carbon taxation is an effective tool to curb CO_2 emissions, muted public support for such environmental policies has so far been a significant political constraint. Fostering social norms might alleviate these political constraints by increasing support for environmental policies—even if they are individually costly.

Our study builds on and contributes to several strands of the literature. First, we contribute to the literature studying the role of social norms in human behavior (see, e.g., Durlauf and Young, 2001; Bowles, 2004; Young, 2008, 2015; Nyborg, Anderies, Dannenberg, Lindahl, Schill, et al., 2016; Nyborg, 2018). We extend this literature and show that individual beliefs about prevalent climate behaviors and norms strongly predict individual willingness to fight climate change. Importantly, we document that Americans vastly *underestimate* the true share of their fellow citizens who try to fight or think that Americans should try to fight global warming. We show that correcting these misperceptions leads to a significant increase in individual willingness to fight climate change and increases individual support for climate-friendly public policies.[1]

Misperceptions of social norms have been documented in settings where social norms are in a phase of transition, giving rise to a phenomenon referred to as '*pluralistic ignorance*' (Allport, 1924; Miller and McFarland, 1987). The majority of a population may privately endorse a norm but incorrectly assume that it is not endorsed by others. This incorrect belief may discourage people from endorsing the norm in public, thereby confirming other people's pessimistic beliefs. For instance, Kuran (1991) argues that a misperception of others' attitudes delayed the collapse of the communist regime in the Soviet Union. More recently, Bursztyn, Rao, Roth, and Yanagizawa-Drott (2020) study the role of misperceived social norms regarding female labor force participation in Saudi Arabia. Our evidence suggests that pluralistic ignorance exists in the context of climate norms and that a low-cost intervention has the potential to significantly alter individual willingness to fight climate change. Thereby, we contribute to recent work which shows that misperceptions about others' behavior, traits, and attitudes are widespread (Bursztyn and Yang, 2021). For instance, research in psychology and political science documents that people tend to underestimate how many of their fellow citizens believe that climate change is real and dangerous (Leviston, Walker, and Morwinski, 2013; Geiger and Swim, 2016;

1. Related to our work are recent studies showing that informational interventions that raise people's awareness about their neighbors' energy consumption or water use causally affect energy or water demand (see, e.g., Allcott (2011), Costa and Kahn (2013), Ferraro and Price (2013), and Jachimowicz, Hauser, O'Brien, Sherman, and Galinsky (2018)). In contrast to these studies, we provide causal evidence that (misperceived) social norms play a role in determining individual willingness to fight climate change and support for public policies. Our study thus also differs from recent correlative analyses that find a positive association between norm perception and environmental behavior (Farrow, Grolleau, and Ibanez, 2017; Valkengoed and Steg, 2019).

Pearson, Schuldt, Romero-Canyas, Ballew, and Larson-Konar, 2018; Mildenberger and Tingley, 2019; Ballew, Rosenthal, Goldberg, Gustafson, Kotcher, et al., 2020).

Moreover, we contribute to the literature examining the relationship between economic preferences and human behavior. Fundamental economic preferences such as time preferences, risk preferences, or prosociality have been shown to predict a wide range of human behaviors (see, e.g., Barsky, Juster, Kimball, and Shapiro, 1997; Dohmen, Falk, Huffman, and Sunde, 2009; Dohmen, Falk, Huffman, Sunde, Schupp, et al., 2011; Falk, Becker, Dohmen, Enke, et al., 2018; Figlio, Giuliano, Özek, and Sapienza, 2019). They have also been shown to predict a set of specific pro-environmental behaviors such as individual willingness to save energy or invest in energy-efficient technology (see, e.g., Newell and Siikamki, 2015; Schleich, Gassmann, Meissner, and Faure, 2019; Fischbacher, Schudy, and Teyssier, 2021; Lades, Laffan, and Weber, 2021). In contrast to these studies, we examine the relationship between economic preferences and individual willingness to fight climate change – as measured through an incentivized donation decision – in a large, representative sample of US adults. The decision to give up money to protect the climate reflects a central trade-off that individuals face when deciding whether to engage in climate-friendly behavior. This allows us to abstract from ancillary factors that are likely to shape specific pro-environmental decisions but are context-specific (e.g., the riskiness of investments in energy-efficient technology).

Finally, we explore the relationship between universal moral values and individual willingness to fight climate change.[2] Recent advances in moral psychology posit that people's moral values can be partitioned into different moral foundations and that holding universal moral values predicts individual behaviors such as voting or support for polices such as environmental protection (Haidt and Joseph, 2004; Haidt, 2012; Graham et al., 2013; Enke, Rodríguez-Padilla, and Zimmermann, 2019; Enke, 2020; Welsch, 2020). We show that universal moral values predict climate preferences over and above what can be predicted by economic preferences such as social preferences. Holding universal moral values might be particularly relevant in the context of climate change, where local behavior has consequences for people around the globe.

3.2 Study Design

To study individual willingness to fight climate change and its behavioral determinants, it is important to obtain a reliable and inter-personally comparable measure of individual willingness to fight climate change as well as detailed information on its potential determinants, such as perceived social norms, fundamental economic preferences, and moral values. To make inferences about the US population, a large

2. See Drews and Bergh (2016), Gifford (2011), or Swim, Clayton, Doherty, Gifford, Howard, et al. (2009) for broad reviews of other determinants of climate behavior and climate policy support.

representative sample is required. Establishing a causal relationship between perceived social norms and climate behavior further requires exogenous variation in the perception of norms. This section explains how we design the sampling approach and survey to meet these requirements.

3.2.1 Sample and survey procedures

We collect survey data from a representative sample of 8,000 study participants in the US. To be eligible to participate in the study, respondents had to reside in the US and be at least 18 years old. The data collection was carried out in two waves. The first wave of data ($N = 2,000$) was collected in March 2021. This wave of data forms the basis for the descriptive analysis presented in this paper, and informs the treatments embedded into wave 2. The second wave of data ($N = 6,000$) was collected in April 2021 and it contains the information experiment that allows us to study the causal relationship between perceived social norms and individual willingness to fight climate change.[3]

We used a stratified sampling approach to ensure that the samples represent the adult US population in terms of gender, age, education, and region. Comparing our samples to data from the American Community Survey 2019, we note that the distribution of demographic characteristics in our samples closely matches the distribution of characteristics in a nationally representative sample (see Appendix Table 3.A.1).

The survey contains several modules. In the following, we explain how we measure individual willingness to fight climate change (Section 3.2.2) and proceed with describing our measures of potential determinants (Section 3.2.3). We then present the information intervention embedded into wave 2 and explain how we elicit posterior beliefs (Section 3.2.4). We also measure individual support for climate policies, political engagement, climate change skepticism and a range of background characteristics (Section 3.2.5). The exact wording of the main survey blocks is provided in Appendix 3.B.

3.2.2 Measuring individual willingness to fight climate change

To measure individual willingness to fight climate change, we use an incentivized donation paradigm. Respondents are asked to divide $450 between themselves and

3. To collect the data, we collaborated with the professional survey company *Pureprofile*, which is frequently used in social science research. All survey participants were part of the company's online panel and participated in the survey online. The online surveys were scripted in the survey software Qualtrics. In both waves, the median time to complete the survey was 18 minutes. Respondents could only participate in one of the two waves. We screen out participants who do not pass an attention check (see Appendix 3.B.1) or speed through the survey with a duration of less than three minutes. Both exclusion criteria are pre-registered.

atmosfair, a charitable organization that fights global warming and offsets CO_2 emissions.[4] The more money that a respondent is willing to donate, the higher their willingness to fight climate change. The measure is quantitative and inter-personally comparable, and it captures the central trade-off underlying most individual-level decisions to fight climate change: mitigating climate change comes at a cost, whether in terms of money, time, or convenience. The amount of $450 was chosen because, by donating the full amount, respondents could offset the annual CO_2 emissions of an average US citizen.[5] We explain this to respondents in order to put their contribution decision into context and render it meaningful and tangible.

Before respondents make their decision, the instructions provide further information on *atmosfair*. Participants are informed that the charity actively contributes to CO_2 mitigation by promoting, developing, and financing renewable energies worldwide. Further information is provided on the charity's annual expenditure dedicated to the fight against global warming ($12 million) as well as its low overhead costs (5%). To minimize rounding, respondents can indicate their responses using a slider ranging from $0 to $450.

The incentive scheme is probabilistic: 25 participants are chosen at random and their decisions are implemented accordingly. The use of high-stake incentives mitigates the problem of experimenter demand effects or social desirability bias that might be present in hypothetical decisions.

3.2.3 Measuring behavioral determinants

Perceived social norms. Social norms are behavioral rules that express the collectively shared understanding of what is typical and morally acceptable behavior. They set the standards of conduct, shape individual behavior, are decentrally enforced, and could thus create a potent momentum either in favor of or against climate action (Bicchieri, 2006; Krupka and Weber, 2013; Nyborg et al., 2016; Bursztyn and Jensen, 2017). We thus hypothesize that individual willingness to fight global warming is determined by individual perceptions of other people's behavior (*'perceived behavior'*) as well as individual perceptions of what other people believe should be done (*'perceived norms'*).[6] Beliefs about the choices that other people make reflect

4. Throughout the survey, we use the term "global warming" instead of the preferred scientific term "climate change" as the former is less likely to be confused with short term or seasonal weather changes or ozone depletion, a misunderstanding that still occasionally arises (Lorenzoni, Leiserowitz, Franca Doria, Poortinga, and Pidgeon, 2006). To avoid confusion, we define global warming as follows at the beginning of the survey: "*Global warming means that the world's average temperature has considerably increased over the past 150 years and may increase more in the future.*" Throughout this text, we use the terms global warming and climate change interchangeably.

5. At the time of the survey, it cost about $28 to offset 1 ton of CO_2 emissions. The World Bank estimates that a typical US resident causes about 16 tons of CO_2 emissions per year.

6. The former are sometimes referred to as descriptive norms or empirical beliefs, while the latter are also sometimes referred to as second-order normative beliefs, injunctive norms, or prescriptive norms (Cialdini, Reno, and Kallgren, 1990; Bicchieri, 2006).

the perceived behavioral standard or norm in a community, which is particularly relevant when people condition their cooperation on the action of others ('conditional cooperation', Fischbacher, Gaechter, and Fehr, 2000). Beliefs about what other people consider appropriate reflect the perceived moral rules or principles in a community. People might have a preference to adhere to the prevalent rules to protect their reputation or self-image (Bursztyn and Jensen, 2017; Falk, 2021).

Before eliciting respondents' perceptions about prevalent social norms, we first ask respondents two questions which allow us to establish prevalent behaviors and endorsement of norms in a representative sample of US adults. Specifically, we ask all respondents about their own behavior and endorsement of the norm to fight global warming: Do they "try to fight global warming" (yes/no)? Do they believe "people in the US should try to fight global warming" (yes/no)?

To measure *perceived* social norms, we then ask all respondents to estimate what proportion of the US population "try to fight global warming" ('perceived behavior') and what proportion think that "people in the US should try to fight global warming" ('perceived norms'). Before making their guesses, respondents are informed that we have gathered survey evidence on whether people try to fight global warming and whether they think that people in the US should try to fight global warming. More specifically, it is explained that we have surveyed a large sample of the US population and that the survey results "represent the views and attitudes of people in the United States". For ease of comprehension, respondents are not asked to estimate proportions but rather estimate the number of people to whom the statement applies out of 100 people we asked:

- *Out of 100 people we asked, how many stated that they try to fight global warming?*
- *Out of 100 people we asked, how many stated that they think that people in the United States should try to fight global warming?*

To determine whether individual perceptions are correct, we can compare participants' guesses with the actual shares of wave 1 respondents answering affirmatively to the questions whether they "try to fight global warming" and whether they think that "people in the US should try to fight global warming". We incentivize the guesses that respondents make to induce and reward careful and accurate responses. In particular, every respondent can earn a $1 bonus if their guess in a randomly-selected belief question differs at most by three from the true value.[7] The resulting measures of perceived behaviors and perceived norms are simple, yet quantitative, incentivized, and inter-personally comparable. Together, they capture the two key facets of social norms that have been identified as key drivers of human behavior in many contexts.

7. The perceived behavior and the perceived norms question are the central but not the only belief questions in the survey. In total, we ask fifteen different belief questions, all of which are incentivized by the reward scheme. The additional belief questions are introduced in Section 3.2.4.

Economic preferences. Economic preferences have been shown to predict a range of important decisions and they are likely to be important determinants of individual willingness to fight climate change. To explore the relationship between economic preferences and the propensity to fight global warming, we obtain detailed individual-level measures of economic preferences following the methodology used in the Global Preferences Survey (Falk, Becker, Dohmen, Enke, et al. (2018) and Falk, Becker, Dohmen, Huffman, et al. (2018)). This experimentally validated survey relies on a range of qualitative and quantitative survey items and allows us to construct preference measures for six fundamental preferences: *patience*, *willingness to take risks*, *altruism*, *trust*, *positive reciprocity*, and *negative reciprocity*. The latter two capture the willingness to reward kind or punish unkind actions, respectively. More information on the survey items and how the composite measures are computed can be found in Appendix 3.C. For ease of interpretation, each preference measure is standardized to have a mean of zero and a standard deviation of one.

Universal moral values. Moral universalism captures the tendency to extend altruistic and moral concerns to individuals who are socially distant (Singer, 2011; Crimston, Bain, Hornsey, and Bastian, 2016; Enke, 2020). Given the global nature of climate change, there are strong reasons to hypothesize that individual willingness to fight global warming is determined by the relative importance of universal versus communal moral values. Moral Foundations Theory (MFT) posits that people's moral concerns can be partitioned into five distinct foundations: care/harm, fairness/reciprocity, in-group/loyalty, authority/respect, and purity/sanctity. "Universal" values – captured by the care/harm and fairness/reciprocity foundations – apply irrespective of the people involved. "Communal" values – captured by the in-group/loyalty and authority/respect foundations – are tied to certain groups or relationships (Haidt and Joseph, 2004; Haidt, 2012; Graham et al., 2013; Enke, 2020). We administer the Moral Foundations Questionnaire (MFQ) to measure the distinct foundations and calculate the relative importance of universal moral values following the approach proposed by Enke (2020). More information on how the standardized measure is constructed can be found in Appendix 3.C.

3.2.4 Shifting perceived social norms

Given the threat posed by global warming, it is important to understand which interventions could increase individual willingness to fight climate change. While it is difficult to alter fundamental human traits such as altruism, patience, or moral values in the short term, beliefs about social norms are likely to be considerably more malleable.[8] As we will show in Section 3.3.2, respondents on average misperceive

8. Economic preferences such as altruism and patience are also malleable, especially during the childhood period, and can be affected through educational interventions in the case of patience (Alan and Ertac, 2018) or through an enriched social environment in the case of altruism (Kosse, Deckers,

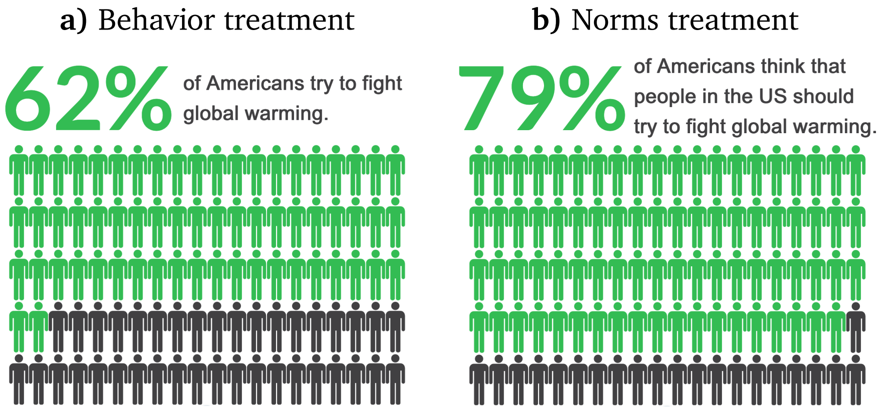

Notes: Panels a and b provide a visual summary of the information provided to participants in the behavior and the norms treatments, respectively. The exact wording of the survey instructions is provided in Appendix 3.B.

Figure 3.2.1. Information treatments in wave 2

the prevalence of social norms in the US. Motivated by this finding, we embed an information experiment into wave 2. The exogenous variation induced by this experiment allows us to study whether the perceived prevalence of social norms causally affects individual willingness to fight global warming.

After eliciting respondents' beliefs about prevalent behaviors and norms, we provide randomly-selected participants with truthful information about the proportion of the US population who (i) "try to fight global warming" ('*behavior treatment*') or (ii) think that "people in the US should try to fight global warming" ('*norms treatment*'). Estimates of both shares are derived from wave 1. More specifically, we randomize respondents in wave 2 into one of three treatments. Appendix Figure 3.A.1 summarizes the structure of the experiment.

1. Behavior treatment In this treatment, respondents are informed about the share of the US population who "try to fight global warming". Respondents are first informed about the fact that "we recently surveyed 2,000 people in the United States and asked them whether they try to fight global warming. Respondents come from all parts of the population and their responses represent the views and attitudes of people in the United States." On the following page, respondents learn that 62% of Americans try to fight global warming. To ensure that participants pay attention, the information is revealed piece by piece, and respondents need to spend a minimum of 5 seconds on the final screen before being able to proceed. A graph on the final screen expresses the information visually, making it salient and tangible (see Figure 3.2.1.a).

Pinger, Schildberg-Hörisch, and Falk, 2019; Rao, 2019). While it is possible that such interventions can lead to an increased willingness to fight climate change, these interventions are more difficult to implement on a larger scale.

2. Norms treatment In an analogous manner, respondents in the norms treatment learn that 79% of Americans think that people in the US should try to fight global warming (see Figure 3.2.1.b).

3. Control No information is provided to participants in the control condition.

Subsequently, we elicit individual willingness to fight climate change with the incentivized donation decision (see Section 3.2.2), which constitutes our main outcome measure. This study design allows us to assess whether providing respondents with accurate information about prevalent behaviors or norms can shift individual climate behavior.

Respondents randomized into the behavior or norms treatment are likely to revise their beliefs about prevalent behaviors or norms in the US. Such a shift in beliefs may lead to a change in individual willingness to fight climate change. Since – as we will show – individuals systematically underestimate the share of Americans trying to fight global warming as well as the share who think that Americans should try to fight global warming, we posit that the information interventions are likely to increase individual willingness to fight climate change. We opt for the dual approach of shifting both perceived behavior and perceived norms, as both are regarded as central drivers of human behavior. However, conceptually, these two entities are closely related. A change in perceived behavior may also lead to a change in perceived norms and vice versa. We explore this question in further detail in Section 3.3.3.

To study belief revisions, we include a post-treatment module in which we elicit posterior beliefs. Respondents are asked to estimate what proportion of the US population engages in a set of concrete climate-friendly behaviors ('perceived behaviors') and what proportion of the US population thinks that one should engage in those behaviors ('perceived norms'). The set of concrete behaviors includes restricting meat consumption, avoiding flights, using environmentally-friendly alternatives to fossil-fueled cars, using green electricity, adapting shopping behavior to the carbon footprint of products, and politically supporting the fight against global warming. Guesses are incentivized using the same reward scheme as described in Section 3.2.3. To determine whether guesses are correct, we compare individual responses to the actual share of wave 1 respondents who report engaging in these behaviors or stating that they think one should engage in those behaviors. For the purpose of the analysis, we compute a perceived behavior index and a perceived norms index by calculating the average across the six climate-friendly behaviors/norms items. We then standardize each index to have a mean of zero and a standard deviation of one among control group respondents. Conceptually, individual perceptions about the prevalence of concrete behaviors/norms are strongly related to the more general behavior/norm of "trying to fight global warming". We can thus use those questions to test for and detect belief revisions without repeating our main questions, thereby mitigating experimenter demand effects and consistency bias in survey responses (Haaland et al., forthcoming).

3.2.5 Additional measures

Climate change skepticism. The public and political debate on climate change has been shaped by a denial of its existence, dangers, or human origin. This phenomenon is particularly relevant in the US where climate change skepticism is widespread and has often formed a key obstacle to effective responses against climate change (Dunlap and McCright, 2011; Leiserowitz, Maibach, Roser-Renouf, Smith, and Dawson, 2013). The subgroup of climate change deniers thus holds particular political relevance, and the survey includes a diverse set of items that allow us to measure respondents' skepticism. We ask respondents to indicate how much trust they have in climate science, whether they think scientists agree that global warming is happening, how worried they are about global warming, whether they think it will harm people in the US, and whether they think that climate change is human-caused (see Appendix 3.B). These questions are asked at the beginning of the survey to ensure that the responses are not affected by the information treatments. We use this information to explore the heterogeneity of treatment effects.

Policy support and political engagement. In addition to eliciting individual willingness to fight climate change, we collect detailed information on the extent to which individuals support different climate polices (e.g., a carbon tax, subsidies for green energy, pollution regulation) and are willing to engage politically (e.g., volunteer time, attend protest, contact government officials). We pose a total of 18 questions adapted from a detailed politics module developed as part of the Climate Change in the American Mind Project (Howe, Mildenberger, Marlon, and Leiserowitz, 2015). Respondents can express their policy support and individual political engagement on a four-point Likert scale (see Appendix 3.B), which we recode in our analysis to ensure that larger values indicate more policy support and political engagement. For ease of interpretation, we aggregate individual items into a policy support index (7 items), a political engagement index (11 items), and a joint index comprising all 18 items. Each index is standardized to have a mean of zero and a standard deviation of one among control group respondents. The questions are posed after the information treatments in wave 2, which allows us to study whether shifting beliefs about prevalent behaviors and norms causally affects policy support and willingness to engage politically.

Background characteristics. We collect detailed information on individual background characteristics. Those include age, gender, education, employment status, household income, the number of children, and whether the respondent thinks of themselves as being closer to the Republican or Democratic party. We use those variables as additional control variables in the analysis.

3.3 Results

3.3.1 Willingness to fight climate change and its determinants

To measure climate preferences, we use an incentivized donation decision in which respondents divide \$450 between themselves and a charitable organization that fights global warming. We use this measure to study how climate preferences are distributed across the population and examine which factors predict those preferences. For the purpose of this descriptive analysis, we focus on survey data collected in wave 1 ($N = 2,000$), which did not contain any treatment manipulation.

Appendix Figure 3.A.2 displays the distribution of individual willingness to fight global warming, as measured through the incentivized donation decision. On average, respondents are willing to donate \$225 of the \$450. There is a considerable degree of heterogeneity across respondents, with 6% donating \$0, 12% donating \$450, and the remaining 82% donating some value in between.

We explore which factors predict individual willingness to fight climate change. For this purpose, we regress the donation amount (in \$) on (i) individual beliefs about prevalent behaviors or norms, (ii) our measures of fundamental economic preferences (i.e., patience, risk-taking, altruism, positive reciprocity, negative reciprocity, and trust), (iii) universal moral values, and (iv) a range of background characteristics. Given that beliefs about prevalent behaviors and norms are conceptually related and highly correlated in our data ($\rho = 0.67$), we estimate two separate regression models, including one belief measure at a time. For the purpose of this analysis, the belief measures are standardized to have a mean zero and a standard deviation of one. The results are reported in columns 1 and 2 of Table 3.3.1, respectively.

First, perceived behaviors and norms are strong predictors of climate preferences. Controlling for the large set of covariates, a one-standard-deviation increase in perceived behavior is associated with a \$12 higher donation amount ($p < 0.001$), while a corresponding increase in perceived norms is associated with a \$14 higher donation ($p < 0.001$).[9] These results are consistent with norm perceptions playing an important role in determining individual willingness to fight global warming. This could, for example, be the case if individuals are 'conditional cooperators' or if they have a preference for complying with existing social norms. Whether or not this relationship can be interpreted as causal is a question we turn to in Section 3.3.3.

Second, the results presented in Table 3.3.1 further reveal that climate donations are fundamentally related to economic preferences. Altruism and positive reciprocity – both of which are facets of prosociality – positively predict the donation amount. The magnitudes of the estimated coefficients are sizeable. For example, a

9. We note that both belief measures have a standard deviation of 22 percentage points. The coefficients can therefore also be interpreted as follows: A 10 percentage point increase in the behavior belief is associated with a \$5.50 higher donation amount, while a corresponding increase in the norms belief is associated with a \$6.50 higher donation amount.

Table 3.3.1. Determinants of climate change behavior

	Donation ($)	
	(1)	(2)
Perceived social norms		
Behavior belief	12.237***	
	(3.154)	
Norms belief		14.500***
		(3.058)
Economic preferences		
Altruism	51.267***	51.734***
	(3.477)	(3.448)
Patience	15.195***	15.192***
	(3.105)	(3.096)
Risk	-1.411	-0.792
	(3.373)	(3.354)
Positive reciprocity	9.571***	7.877**
	(3.239)	(3.258)
Negative reciprocity	-3.338	-2.540
	(3.214)	(3.185)
Trust	1.071	0.831
	(3.233)	(3.203)
Moral foundations		
Relative universalism	23.772***	23.420***
	(3.301)	(3.290)
Sociodemographics		
Democrat	45.143***	44.160***
	(6.241)	(6.246)
Age	0.685	0.702
	(1.035)	(1.034)
Age (squared)	-0.007	-0.006
	(0.011)	(0.011)
Female	16.943***	16.520***
	(6.367)	(6.331)
Log income	9.965***	9.895***
	(3.741)	(3.726)
College degree	-15.320**	-15.953**
	(6.522)	(6.504)
Employed	8.453	8.868
	(6.661)	(6.638)
Parent	4.659	4.695
	(6.498)	(6.478)
R^2	0.281	0.284
N	1,975	1,975
Mean of dep. var.	225.21	225.21

Notes: This table shows OLS regression estimates using respondents from wave 1, where the dependent variable is the amount donated to the charitable organization that fights global warming. Perceived social norms, economic preferences, and universal moral values are standardized. “Democrat”, “Female”, “College degree”, “Employed” and “Parent” are binary indicator variables. “Log income” is coded as the log of the income bracket’s midpoint. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

one-standard-deviation increase in altruism is associated with a $52 higher donation amount. Similarly, patience positively predicts donation decisions. These patterns are plausible given that climate action benefits other people around the world as well as future generations. We find no statistically significant associations between climate preferences and risk preferences, negative reciprocity, or trust.

Third, we find a strong positive association between universal moral values and climate preferences. A one-standard-deviation increase in relative universalism – namely the extent to which individuals endorse universal moral values that apply equally to all humans rather than communal or ingroup-restricted values – is associated with a $23 higher donation amount. Climate change is a global problem and individuals whose moral values apply irrespective of the people involved are more likely to make larger donations, presumably because they are more likely to take the welfare of other people outside of their community into account.

Finally, demographic characteristics also significantly predict individual willingness to fight climate change. Democrats on average contribute about $45 more than Republicans, female respondents about $16 more, and household income is also positively associated with the donation amount. However, higher education negatively predicts climate donations. Further analyses reveal that this effect is entirely driven by Republicans among whom a college education is associated with a $27 lower donation amount (see Appendix Table 3.A.2).[10]

Taken together, the results suggest that perceived social norms, economic preferences, and universal moral values are likely to shape individual willingness to fight climate change. Since climate action is commonly conceived as a global and intergenerational cooperation problem, it seems plausible that a higher willingness to fight climate change requires some degree of prosociality, patience, and universal moral values. Beliefs about prevalent behaviors and norms are also likely to be key determinants of individual willingness to fight global warming if individuals act as ‘conditional cooperators’ or have a preference to comply with existing social norms.

3.3.2 Misperceived social norms

Having established which factors are predictive of individual willingness to fight climate change, we now explore the distribution of beliefs about behaviors and norms in more detail. Given that these beliefs are potentially malleable, it holds particular importance to establish whether there are systematic misperceptions of prevalent behaviors and norms. For the purpose of this analysis, we again rely on the survey data collected in wave 1.

10. We are not the first to document a negative education gradient among Republicans (Hamilton, 2011; Newport and Dugan, 2015). It has been hypothesized that highly-educated individuals are cognitively better equipped to rationalize and internalize the views of their cultural community, which for Republicans might correspond to climate change skepticism (Kahan, Peters, Wittlin, Slovic, Ouellette, et al., 2012; but see Van Der Linden, Maibach, Cook, Leiserowitz, Ranney, et al., 2017).

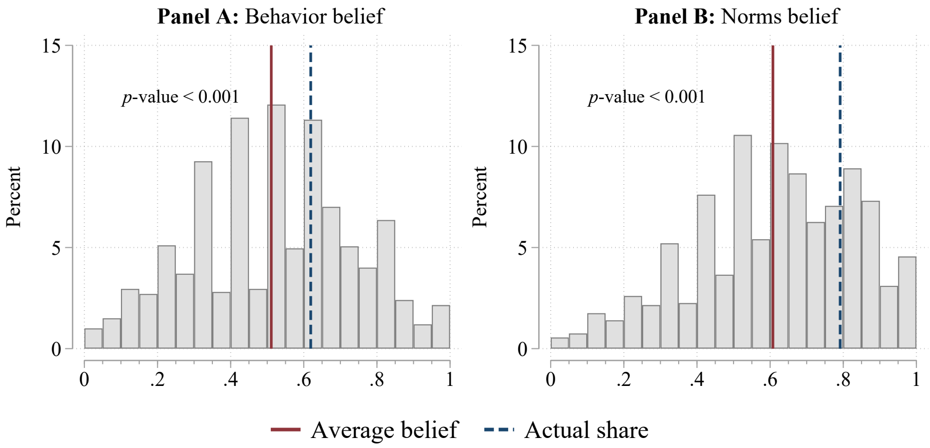

Notes: This figure shows the distribution of perceived social norms in wave 1. Panel A shows the distribution of people's beliefs about the share of Americans who say that they try to fight global warming. Panel B shows the distribution of people's beliefs about the share of Americans who say that one *should* fight global warming. Each panel indicates the average belief across respondents (solid red) as well as the actual shares (dashed blue) as vertical lines.

Figure 3.3.1. Perceived social norms: fight global warming

Figure 3.3.1 depicts the distribution of perceived social norms. Panel A displays perceived behavior, i.e., the distribution of individual beliefs about the share of the US population that tries to fight global warming. Panel B displays perceived norms, i.e., the distribution of beliefs about the share of Americans who think that people in the US should try to fight global warming. The average belief is indicated by a vertical red line, whereas the actual share is marked by a dotted blue line.

Figure 3.3.1 reveals a considerable degree of heterogeneity in individual beliefs. Both panels further reveal that respondents vastly misperceive the prevalence of climate-friendly behaviors and norms among their fellow citizens. On average, respondents believe that 51% of Americans try to fight global warming, while the actual share is 62% (p-value < 0.001). The majority of participants – namely 67% – underestimate how prevalent climate-friendly behavior is in the US. Similarly, respondents on average believe that 61% of Americans think that people in the US should try to fight global warming, while the actual share is 79% (p-value < 0.001). Again, most participants (76%) underestimate this share.[11] We find larger misperceptions among respondents who are older, have a lower income, have a lower education, or are Republicans (see Appendix Table 3.A.4).

Taken together, while the majority of Americans try to fight global warming and a vast majority agrees that people in the US should try to fight global warming, most

11. We also elicit beliefs about concrete climate change behaviors, e.g., restricting meat consumption, avoiding flights and cars, or consuming only green electricity. These measures are highly correlated with the abstract measure (see Appendix Table 3.A.5). Moreover, Appendix Figures 3.A.4 and 3.A.5 show that we document similar norm misperceptions for these concrete behaviors.

Americans underestimate the degree to which other Americans engage in climate-friendly behaviors and share those normative views. This underestimation of climate norms is likely to hamper individual willingness to fight climate change.

3.3.3 Correcting misperceived social norms

As established in the previous sections, beliefs about prevalent behaviors and norms strongly predict individual willingness to fight climate change. At the same time, there are systematic misperceptions of the actual share of Americans fighting or thinking that one should fight climate change. Can information interventions that inform respondents about the true shares can affect individual willingness to fight climate change? The information experiment embedded in wave 2 allows us to study this question. Respondents are randomized into (i) a 'behavior treatment', in which they are informed that 62% of Americans try to fight global warming, (ii) a 'norms treatment', in which they are informed that 79% of Americans think that people in the US should try to fight global warming, or a (iii) a control group. Appendix Table 3.A.3 presents the balancing of characteristics across the three groups. We cannot reject the null hypothesis that the three groups differ in terms of observable characteristics and conclude that the randomization was successful. Appendix Figure 3.A.3 displays the wedge between wave 2 respondents' beliefs about prevalent behaviors and norms and the actual shares. As can be seen from both figures, wave 2 participants also vastly underestimate the true shares, providing us with an ideal opportunity to exogenously correct inaccurate perceptions. The average gap between the perceived and actual shares is 10 percentage points in the case of perceived behaviors and 17 percentage points in the case of perceived norms.

To estimate the causal impact of the information treatments, we regress willingness to fight climate change – as measured through the incentivized donation decision (in \$) – on treatment indicators and a set of control variables.[12] The results are reported in column 1 of Table 3.3.2 and reveal that the impacts of the information treatments are sizeable and highly statistically significant. Being informed about the true share of Americans who try to fight global warming leads to a \$12 increase in donations (p-value = 0.012), while being informed about the true share of Americans who think that people in the US should try to fight global warming increases donations by \$16 (p-value < 0.001). The effects correspond to a relative increase of 4.7% and 6.3%, respectively. While the point estimate of the coefficient for the norms treatment is somewhat larger than the point estimate of the coefficient for the behavior treatment, we note that the two are not significantly different from each other (p-value = 0.39). Given that not all respondents misperceive prevalent

12. The set of control variables includes controls for gender (indicator), age (continuous), log income, college degree (indicator), employment (indicator), party affiliation (indicator), and census region (three indicators). Appendix Table 3.A.6 presents results of the regressions without control variables. The estimated coefficients are very similar in magnitude and significance.

Table 3.3.2. Treatment effects on climate donations and posterior beliefs

	(1) Donation ($)	(2) Behavior belief (post.)	(3) Norms belief (post.)
Behavior treatment	11.725**	0.279***	0.235***
	(4.675)	(0.030)	(0.030)
Norms treatment	15.674***	0.370***	0.350***
	(4.701)	(0.031)	(0.030)
N	5,991	5,988	5,976
Control group mean	249.31	0	0
z-scored	No	Yes	Yes
Controls	Yes	Yes	Yes

Notes: This table shows OLS regression estimates using respondents from wave 2. The dependent variable is the donation to the climate charity (in $). It is regressed on binary indicators that take the value of 1 for respondents in the behavior treatment and norms treatment, respectively. "Behavior belief" is an index of six post-treatment beliefs about the share of Americans engaging in concrete climate-friendly behaviors to fight global warming. "Norms belief" is an index of six post-treatment beliefs about the share of Americans who say that one should engage in concrete climate-friendly behaviors to fight global warming. Both indices are standardized to have a mean of zero and a standard deviation of one in the control group. All regressions include controls for gender (indicator), age (continuous), log income, college degree (indicator), employment (indicator), party affiliation (indicator), and census region (three indicators). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

behaviors and norms at the baseline and some respondents might not fully revise their beliefs in light of the information provided, both effect sizes suggest a powerful impact of perceived social norms on individual willingness to fight climate change.[13]

Using the posterior norm perception module, we provide evidence that the treatments indeed shift posterior beliefs in the way that one would expect. To study belief revisions, we regress the posterior beliefs about concrete climate-friendly behaviors and norms on the treatment indicators and the same set of control variables. As explained in Section 3.2.4, the set of concrete behaviors includes different actions such as reducing meat consumption or avoiding flights. The two posterior belief indices are standardized, and the results are reported in columns 2 and 3 of Table 3.3.2, respectively. Both information treatments successfully shift beliefs, which are revised upwards by 0.24 to 0.37 standard deviations. We also observe spill-over effects. Information about prevalent behavior also shifts beliefs about prevalent norms and vice versa. As remarked earlier, the treatments should not be interpreted as separate manipulations of orthogonal concepts but rather as statistically independent

13. We can derive the treatment effect per standardized change in beliefs under the assumption that respondents fully update their beliefs to the information provided, which implies an average belief increase of 0.47 standard deviations in the behavior treatment and 0.82 standard deviations in the norms treatment. The behavior treatment thus has a $24.8 effect and the norms treatment a $19.2 effect on climate donations per standardized belief change. Both figures likely underestimate the true effect because most respondents presumably only partially update their beliefs.

Table 3.3.3. Treatment effect heterogeneity: Prior above/below actual share

	Dependent variable: Donation ($)			
	Prior < actual share		Prior ≥ actual share	
	(1)	(2)	(3)	(4)
Behavior treatment	14.931**		5.231	
	(5.875)		(7.701)	
Norms treatment		19.111***		4.747
		(5.387)		(9.623)
N	2,579	3,054	1,399	946
Control group mean	243.09	241.67	260.69	273.71
Controls	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression estimates using respondents from wave 2. The dependent variable is the donation to the climate charity (in $). It is regressed on binary indicators that take the value of 1 for respondents in the behavior treatment and norms treatment, respectively. We run separate analyses for respondents with prior norm perceptions strictly below the actual share (columns 1-2) and equal to or above the actual share (columns 3-4). We consider beliefs about others' behavior in the behavior treatment and beliefs about others' norms in the norms treatment. Given that the actual shares are different for the two beliefs, we do not pool all three treatment groups in this analysis. Instead, we only use respondents in the control condition and the behavior treatment in the analysis presented in columns 1 and 3, and only use respondents in the control condition and the norms treatment in the analysis presented in columns 2 and 4. All regressions include the set of controls described in Table 3.3.2. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

yet conceptually-related treatments with a common effect: they both strengthen perceived social norms.

Treatment effect heterogeneity by prior. We explore heterogeneity in treatment effects across different subgroups. First, we examine whether the treatments are more effective for respondents whose priors are below the actual shares. Table 3.3.3 separately displays the treatment effects for respondents whose prior beliefs are below the true shares (Panel A) and those whose prior beliefs are equal to or above the true shares (Panel B). As can be seen from this table, the positive treatment effects that we document for the full sample are almost entirely driven by those individuals whose priors are below the actual shares. Among them, the behavior treatment increases donations by $15 (p-value = 0.011), whereas the norms treatment increases donations by $19 (p-value < 0.001). Reassuringly, we do not observe a back-firing effect. For respondents whose priors are equal to or above the actual shares, the estimated coefficients are positive albeit smaller in magnitude and insignificant.[14]

14. Appendix Figure 3.A.6 displays non-parametric estimates of the moderating role of pre-treatment beliefs for our information treatments (Xu, Hainmueller, Mummolo, and Liu, 2017; Hainmueller, Mummolo, and Xu, 2019). As can be seen from this figure, the effects of the behavior and the norms treatment are stronger among respondents with low pre-treatment beliefs. Moreover, both treatments have a weakly positive effect across the whole belief distribution.

However, we note that we cannot reject the null hypothesis that the treatment effect coefficients are the same for both subgroups.

Treatment effect heterogeneity by climate change skepticism. Next, we explore whether the information treatments lead to a stronger increase in individual willingness to fight climate change for respondents who are skeptical about the existence and threat of human-caused climate change. From a policy perspective, this subset of the population is particularly relevant as it is typically difficult to reach and convince that climate change matters.

Figure 3.3.2 compares the treatment effects across respondents who express skepticism about climate change and those who do not. The sample is split based on five indicators that capture different facets of climate change skepticism: having low trust in climate science, believing that the presence of climate change is still scientifically debated, not being worried about climate change, not perceiving it as a threat for the US, and believing that climate change is mainly the result of natural causes. For all indicators and both treatments, we observe that the point estimates of the treatment coefficients are larger in magnitude for climate change deniers. In the behavior treatment, most coefficients are also statistically different from each other across the two subgroups. For example, the behavior treatment increases donations by $24 for those who report not being worried about global warming and by $39 for those who do not believe that climate change is human-caused. By contrast, we do not find a statistically significant impact of the behavior treatment for respondents who do report being worried or who do believe that climate change is human-caused. These differences in effect sizes are statistically significant at the 5% level (see also Table 3.A.7). In the norms treatment, the differences are more muted.

Climate change deniers tend to have more pessimistic prior beliefs about the prevalence of climate norms in the US. However, we observe largely identical results even if we control for treatment heterogeneity by priors (see Table 3.A.8). Thus, the same information appears to have differential informational value for climate change deniers – even conditional on the same prior belief. Climate change deniers do not only have more scope to adjust their behavior. They might also be surprised to learn that their views are in fact minority views and that the majority of their fellow citizens does take climate change seriously, as indicated by the large share of Americans who take action against it or think that this should be done.[15]

Treatment effects on policy support and political engagement. Do the positive treatment effects of the information treatments also carry over to the political

15. It is unlikely that the much weaker treatment effect among respondents who believe in and are concerned about climate change can be attributed to a “ceiling effect”. In the control treatment, the large majority of these climate change “believers” (about 73% to 75% depending on the question) can still increase their donation by at least $25.

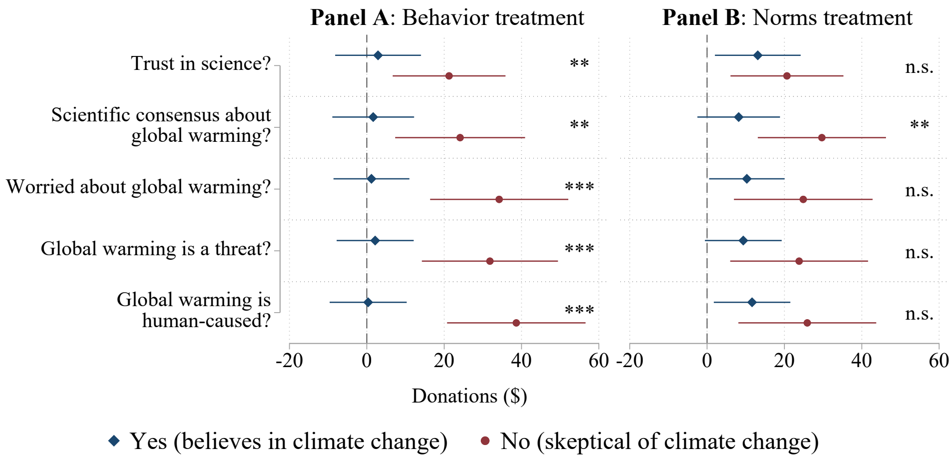

Notes: This figure shows OLS estimates of the treatment effects of the behavior (Panel A) and the norms treatment (Panel B) on donations (in $) in different subsamples. We use respondents from wave 2 and include the set of controls described in Table 3.3.2. 95% confidence intervals are shown. Each panel shows treatment effects among respondents who are skeptical of climate change ("No") and those who believe in climate change ("Yes"), where we use disagreement with different statements as a proxy for skepticism: "Trust in science" means that the respondent trusts climate scientists "a lot" or "a great deal" (on a five-point Likert scale). "Scientific consensus about global warming" means that the respondent thinks that most scientists think that global warming is happening. "Worried about global warming" means that the respondent is "somewhat worried' or "very worried" about global warming (on a four-point Likert scale). "Global warming is a threat" means that the respondent thinks that global warming will do "a moderate amount" or "a great deal" of harm (on a four-point Likert scale). "Global warming is human-caused" means that the respondent thinks that global warming is caused by human activities. For each sample split, we indicate the level of significance of a test of equality of coefficients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, n.s. $p \geq 0.10$.

Figure 3.3.2. Treatment effect heterogeneity by climate change skepticism

domain? To study this question, we collect post-treatment information on policy support and political engagement (see Section 3.2.5). Columns 1 and 2 of Table 3.3.4 present the estimated treatment effects on the standardized indices of support for climate policies and willingness to engage in political actions. Column 3 presents the results for the standardized, joint index. We find that both treatments significantly increase support for climate policies. The behavior treatment significantly increases policy support by 0.09 standard deviations, while the norms treatment significantly increases policy support by 0.07 standard deviations. The estimated coefficients are positive albeit insignificant when we consider willingness to engage in political actions as the outcome. When we use the joint index as the outcome, we find that the behavior treatment significantly increases the index by 0.06 standard deviations, while the norms treatment has an insignificant positive effect of 0.03.

Consistent with the results reported above, we also find that the estimated impacts of the treatments on policy support and political engagement tend to be stronger for the subgroup of climate change deniers. Appendix Figure 3.A.7 shows that both the behavior and the norms treatment significantly increase individual sup-

Table 3.3.4. Treatment effects on support for policies and actions to fight global warming

	(1) Policies	(2) Actions	(3) All
Behavior treatment	0.088***	0.039	0.061**
	(0.026)	(0.027)	(0.026)
Norms treatment	0.066**	0.012	0.034
	(0.026)	(0.027)	(0.026)
N	5,999	5,994	5,993
z-scored	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Notes: This table shows OLS regression estimates using respondents from wave 2. Dependent variables: "Policies" is an index measuring individual support for policies to fight climate change (7 items). "Actions" is an index measuring political engagement in different types of political activities (11 items). "All" is a joint index comprising all 18 items. All indices are constructed by taking the sum of all positively coded items and standardizing the sum to have a mean of zero and a standard deviation of one in the control group. The indices are regressed on binary indicators that take the value of 1 for respondents in the behavior treatment and norms treatment, respectively. All regressions include the set of controls described in Table 3.3.2. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

port for policies to flight global warming by 10 to 20 percent of a standard deviation among climate change deniers. By contrast, our information treatments have hardly any impact on policy support among respondents who believe in climate change.

Taken together, we conclude that providing people with accurate information not only has the potential to increase individual willingness to fight climate change – especially among climate change deniers – but that it can also increase individual support for climate policies.

3.4 Discussion

We document that fundamental human traits such as altruism, positive reciprocity, and moral universalism are strong predictors of individual willingness to fight climate change. This finding could prove fruitful in understanding the frequently observed cultural and political dissent on climate change (Dunlap, McCright, and Yarosh, 2016; Hornsey, Harris, and Fielding, 2018). Indeed, in our data, economic preferences and universalism together explain about 40% of the large $74 baseline donation gap between Republicans and Democrats (see Appendix Table 3.A.9). Likewise, they explain 25% of the gap in policy preferences. Both results suggest that the political divide on climate change can be partially attributed to deeply entrenched human traits. The important role of prosociality further illustrates that many individuals care about the well-being of others and therefore seem to partially internalize the positive externalities of climate action. The traditional economic model of purely self-interested agents facing an insurmountable collective action problem thus underestimates the scope for climate action. Indeed, our survey documents that many Americans are actually willing to act against global warming. 62% of Americans try to fight global warming, and 79% think that this should be done. Moreover, many respondents are willing to give up money to support the work of a climate charity.

Our finding that Americans vastly underestimate the prevalence of climate norms in the US holds particular political relevance. We show both correlationally and causally that perceived social norms are a key driver of individual willingness to fight climate change. The fact that climate norms are commonly underestimated in the US can thus form a dangerous obstacle to climate action. It could trap Americans in an equilibrium with low climate engagement: Individuals are discouraged by the (mis)perceived lack of support, and they abstain from taking actions themselves, which sustains the pessimistic beliefs held by others – a phenomenon that has been dubbed pluralistic ignorance (Allport, 1924; Miller and McFarland, 1987; Bursztyn, Rao, et al., 2020).

However, this diagnosis also implies a unique opportunity to promote and accelerate climate-friendly norms and behavior. We show that a simple, easily scalable, and cost-effective intervention – namely informing respondents about the actual prevalence of climate norms in the US – corrects these misperceptions and encourages climate-friendly behavior. Importantly, we find that this intervention is particularly effective for climate change deniers, namely the group of people who are commonly difficult to reach, but crucial for building up a broad alliance against climate change. Moreover, convincing those who remain skeptical of human-caused climate change is likely to have particularly high returns if these individuals still have ample scope to make their behavior more climate-friendly.

Arguably, the effect of a single, minimalist message as embodied in our information treatments is likely to dissipate with time. However, large-scale information campaigns that repeatedly announce and effectively communicate the actual prevalence of climate norms could correct existing misperceptions and permanently foster climate norms (Bicchieri, 2017). They could trigger a positive feedback loop where learning about the existing support of climate norms encourages Americans to take visible action against climate change, which encourages others to follow suit.

3.5 Conclusion

In this paper, we study the behavioral determinants of individual willingness to fight climate change in a large-scale, representative survey with 8,000 US adults. In a first step, we document that fundamental human traits – namely patience, altruism, positive reciprocity, and moral universalism – are strongly correlated with individual willingness to fight climate change, as measured in a donation decision. Beliefs about the climate behavior and norms of others also matter: Individuals who perceive stronger climate norms are willing to give up more money to support the climate charity. In a second step, we zoom in on perceived social norms, as they are malleable in the short term and can create a potent momentum either in favor of or against climate action. We find that Americans strongly underestimate the support of climate norms in the US. An information experiment shows that informing respondents about the true prevalence of climate norms in the US corrects these misperceptions and increases climate donations.

The widely-observed underestimation of climate norms in the US can form a dangerous obstacle to climate action, whereby moving forward it will be crucial to correct these misperceptions. Our results thus suggest that social norms should play a pivotal role in the policy response to climate change. Policies that foster social norms should complement formal regulations such as carbon taxation. Finally, we hope that the study also showcases an important role that economic and social science research will have to play in the warming years ahead. Its key responsibilities will include monitoring the perception of climate norms, detecting misperceptions early, and exploring how they can effectively be corrected.

References

Alan, Sule, and Seda Ertac. 2018. "Fostering Patience in the Classroom: Results from Randomized Educational Intervention." *Journal of Political Economy* 126 (5): 1865–911. [214]

Allcott, Hunt. 2011. "Social norms and energy conservation." *Journal of Public Economics* 95 (9-10): 1082–95. [209]

Allport, Floyd Henry. 1924. *Social Psychology*. Boston: Houghton Mifflin. [209, 228]

Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk. 2021. "Fighting Climate Change: The Role of Norms, Preferences, and Moral Values." *CEPR Discussion Paper No. DP16343*, 59. URL: https://ssrn.com/abstract=3886831.

Auffhammer, Maximilian. 2018. "Quantifying economic damages from climate change." *Journal of Economic Perspectives* 32 (4): 33–52. [206]

Ballew, Matthew T., Seth A. Rosenthal, Matthew H. Goldberg, Abel Gustafson, John E. Kotcher, Edward W. Maibach, and Anthony Leiserowitz. 2020. "Beliefs about others' global warming beliefs: The role of party affiliation and opinion deviance." *Journal of Environmental Psychology* 70: 101466. [210]

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preferences Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* 112 (2): 537–79. [210]

Bicchieri, Cristina. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press. [212]

Bicchieri, Cristina. 2017. *Norms in the Wild: How to Diagnose, Measure and Change Social Norms*. New York: Oxford University Press. [228]

Bowles, Samuel. 2004. *Microeconomics: Behavior, Institutions, and Evolution*. Princeton: Princeton University Press. [209]

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global non-linear effect of temperature on economic production." *Nature* 527 (7577): 235–39. [206]

Bursztyn, Leonardo, and Robert Jensen. 2017. "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure." *Annual Review of Economics* 9: 131–53. [212, 213]

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." *SSRN Electronic Journal*, DOI: 10.2139/ssrn.3580487. [209, 228]

Bursztyn, Leonardo, and David Y. Yang. 2021. "Misperceptions about Others." *Working Paper*, [209]

Carleton, Tamma A., and Solomon M. Hsiang. 2016. "Social and economic impacts of climate." *Science* 353 (6304): [206]

Cialdini, Robert B., Raymond R. Reno, and Carl a. Kallgren. 1990. "A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places." *Journal of Personality and Social Psychology* 58 (6): 1015–26. DOI: 10.1037/0022-3514.58.6.1015. [212]

Costa, Dora L., and Matthew E. Kahn. 2013. "Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment." *Journal of the European Economic Association* 11 (3): 680–702. [209]

Crimston, Daniel, Paul G. Bain, Matthew J. Hornsey, and Brock Bastian. 2016. "Moral expansiveness: Examining variability in the extension of the moral world." *Journal of Personality and Social Psychology* 111 (4): 636–53. [214]

Diffenbaugh, Noah S., and Marshall Burke. 2019. "Global warming has increased global economic inequality." *Proceedings of the National Academy of Sciences of the United States of America* 116 (20): 9808–13. [206]

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner. 2011. "Individual risk attitudes: Measurement, determinants, and behavioral consequences." *Journal of the European Economic Association* 9 (3): 522–50. [210]

Dohmen, Thomas J, Armin Falk, David Huffman, and Uwe Sunde. 2009. "Homo Reciprocans: Survey Evidence on Prevalence, Behaviour and Success." *Economic Journal* 119 (536): 592–612. DOI: 10.1111/j.1468-0297.2008.02242.x. [210]

Drews, Stefan, and Jeroen C.J.M. van den Bergh. 2016. "What explains public support for climate policies? A review of empirical and experimental studies." *Climate Policy* 16 (7): 855–76. DOI: 10.1080/14693062.2015.1058240. [210]

Dunlap, Riley E., and Aaron M. McCright. 2011. "Organized Climate Change Denial." In *The Oxford Handbook of Climate Change and Society*. Edited by John S. Dryzek, Richard B. Norgaard, and David Schlosberg. Oxford: Oxford University Press. [217]

Dunlap, Riley E., Aaron M. McCright, and Jerrod H. Yarosh. 2016. "The Political Divide on Climate Change: Partisan Polarization Widens in the U.S." *Environment: Science and Policy for Sustainable Development* 58 (5): 4–23. [207, 227]

Durlauf, Steven N., and H. Peyton Young. 2001. *Social Dynamics*. Cambridge: MIT Press. [209]

Enke, Benjamin. 2020. "Moral values and voting." *Journal of Political Economy* 128 (10): 3679–729. [207, 210, 214, 261]

Enke, Benjamin, Ricardo Rodríguez-Padilla, and Florian Zimmermann. 2019. "Moral Universalism and the Structure of Ideology." *Working Paper*, [210]

Falk, Armin. 2021. "Facing yourself - A note on self-image." *Journal of Economic Behavior & Organization* 186: 724–34. [213]

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *Quarterly Journal of Economics* 133 (4): 1645–92. [206, 210, 214, 259]

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. 2018. "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences." *Working Paper*, [206, 214]

Farrow, Katherine, Gilles Grolleau, and Lisette Ibanez. 2017. "Social Norms and Pro-environmental Behavior: A Review of the Evidence." *Ecological Economics* 140: 1–13. [209]

Ferraro, Paul J., and Michael K. Price. 2013. "Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment." *Review of Economics and Statistics* 95 (1): 64–73. [209]

Figlio, David, Paola Giuliano, Umut Özek, and Paola Sapienza. 2019. "Long-Term Orientation and Educational Performance." *American Economic Journal: Economic Policy* 11 (4): 272–309. [210]

Fischbacher, Urs, Simon Gaechter, and Ernst Fehr. 2000. "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." *SSRN Electronic Journal* 71 (3): 397–404. DOI: 10.2139/ssrn.203288. [213]

Fischbacher, Urs, Simeon Schudy, and Sabrina Teyssier. 2021. "Heterogeneous preferences and investments in energy saving measures." *Resource and Energy Economics* 63: 101202. [210]

Geiger, Nathaniel, and Janet Swim. 2016. "Climate of Silence: Pluralistic Ignorance as a Barrier to Climate Change Discussion." *Journal of Environmental Psychology* 47: 79–90. [209]

Gifford, Robert. 2011. “The Dragons of Inaction: Psychological Barriers that Limit Climate Change Mitigation and Adaptation.” *American Psychologist* 66 (4): 290–302. [210]

Graham, Jesse, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean Wojcik, and Peter Ditto. 2013. “Moral Foundations Theory: The pragmatic validity of moral pluralism.” *Advances in Experimental Social Psychology* 47: 55–130. [207, 210, 214]

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2021. “Designing Information Provision Experiments.” *Journal of Economic Literature*, (forthcoming):

Haidt, Jonathan. 2012. *The righteous mind: Why good people are divided by politics and religion*. New York: Vintage. [207, 210, 214]

Haidt, Jonathan, and Craig Joseph. 2004. “Intuitive ethics: How innately prepared intuitions generate culturally variable virtues.” *Daedalus* 133 (4): 55–66. [207, 210, 214]

Hainmueller, Jens, Jonathan Mummolo, and Yiqing Xu. 2019. “How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice.” *Political Analysis* 27 (2): 163–92. [224, 248]

Hamilton, Lawrence C. 2011. “Education, politics and opinions about climate change evidence for interaction effects.” *Climatic Change* 104 (2): 231–42. [220]

Hornsey, Matthew J., Emily A. Harris, and Kelly S. Fielding. 2018. “Relationships among conspiratorial beliefs, conservatism and climate scepticism across nations.” *Nature Climate Change* 8 (7): 614–20. DOI: 10.1038/s41558-018-0157-2. [207, 227]

Howe, Peter D., Matto Mildenberger, Jennifer R. Marlon, and Anthony Leiserowitz. 2015. “Geographic variation in opinions on climate change at a state and local scales in the USA.” *Nature Climate Change* 5 (6): 596–603. [217, 257]

IPCC. 2014. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Edited by Core Writing Team, R. K. Pauchari, and L. A. Meyer. [206]

Jachimowicz, Jon M., Oliver P. Hauser, Julia D. O’Brien, Erin Sherman, and Adam D. Galinsky. 2018. “The critical role of second-order normative beliefs in predicting energy conservation.” *Nature Human Behaviour* 2 (10): 757–64. [209]

Kahan, Dan M., Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel. 2012. “The polarizing impact of science literacy and numeracy on perceived climate change risks.” *Nature Climate Change* 2 (10): 732–35. [220]

Kosse, Fabian, Thomas Deckers, Pia Pinger, Hannah Schildberg-Hörisch, and Armin Falk. 2019. “The Formation of Prosociality: Causal Evidence on the Role of Social Environment.” *Journal of Political Economy* 128 (2): 434–67. [214]

Krupka, Erin L., and Roberto A. Weber. 2013. “Identifying social norms using coordination games: Why does dictator game sharing vary?” *Journal of the European Economic Association* 11 (3): 495–524. DOI: 10.1111/jeea.12006. [212]

Kuran, Timur. 1991. “The East European Revolution of 1989: Is it Surprising that We Were Surprised?” *American Economic Review* 81 (2): 121–25. [209]

Lades, Leonhard K., Kate Laffan, and Till O. Weber. 2021. “Do economic preferences predict pro-environmental behaviour?” *Ecological Economics* 183: 106977. [210]

Leiserowitz, Anthony A., Edward W. Maibach, Connie Roser-Renouf, Nicholas Smith, and Erica Dawson. 2013. “Climategate, Public Opinion, and the Loss of Trust.” *American Behavioral Scientist* 57 (6): 818–37. [217]

Leviston, Z., I. Walker, and S. Morwinski. 2013. “Your opinion on climate change might not be as common as you think.” *Nature Climate Change* 3 (4): 334–37. [209]

Lorenzoni, Irene, Anthony Leiserowitz, Miguel de Franca Doria, Wouter Poortinga, and Nick F. Pidgeon. 2006. "Cross-national comparisons of image associations with "global warming" and "climate change" among laypeople in the United States of America and Great Britain." *Journal of Risk Research* 9 (03): 265–81. [212]

Mildenberger, Matto, and Dustin Tingley. 2019. "Beliefs about Climate Beliefs: The Importance of Second-Order Opinions for Climate Politics." *British Journal of Political Science* 49 (4): 1279–307. [210]

Miller, Dale T., and Cathy McFarland. 1987. "Pluralistic Ignorance: When Similarity is Interpreted as Dissimilarity." *Journal of Personality and Social Psychology* 53 (2): 298–305. [209, 228]

Newell, Richard G., and Juha Siikamki. 2015. "Individual time preferences and energy efficiency." *American Economic Review* 105 (5): 196–200. [210]

Newport, Frank, and Andrew Dugan. 2015. "College-Educated Republicans Most Skeptical of Global Warming." *Gallup Report* March 26: [220]

Nyborg, Karine. 2018. "Social Norms and the Environment." *Annual Review of Resource Economics* 10: 405–23. [209]

Nyborg, Karine, John M. Anderies, Astrid Dannenberg, Therese Lindahl, Caroline Schill, Maja Schlüter, W. Neil Adger, Kenneth J. Arrow, Scott Barrett, Stephen Carpenter, F. Stuart Chapin, Anne Sophie Crépin, Gretchen Daily, Paul Ehrlich, Carl Folke, Wander Jager, Nils Kautsky, Simon A. Levin, Ole Jacob Madsen, Stephen Polasky, Marten Scheffer, Brian Walker, Elke U. Weber, James Wilen, Anastasios Xepapadeas, and Aart De Zeeuw. 2016. "Social norms as solutions." *Science* 354 (6308): 42–43. [209, 212]

Pearson, Adam R., Jonathon P. Schuldt, Rainer Romero-Canyas, Matthew T. Ballew, and Dylan Larson-Konar. 2018. "Diverse segments of the US public underestimate the environmental concerns of minority and low-income Americans." *Proceedings of the National Academy of Sciences* 115 (49): 12429–34. [209]

Rao, Gautam. 2019. "Familiarity Does Not Breed Contempt: Generosity, Discrimination, and Diversity in Delhi Schools." *American Economic Review* 109 (3): 774–809. [215]

Schleich, Joachim, Xavier Gassmann, Thomas Meissner, and Corinne Faure. 2019. "A large-scale test of the effects of time discounting, risk aversion, loss aversion, and present bias on household adoption of energy-efficient technologies." *Energy Economics* 80: 377–93. [210]

Singer, Peter. 2011. *The Expanding Circle: Ethics, Evolution, and Moral Progress*. Princeton University Press. [214]

Swim, Janet, Susan Clayton, Thomas Doherty, Robert Gifford, George Howard, Joseph Reser, Paul Stern, and Elke Weber. 2009. "Psychology & global climate change: Addressing a multifaceted phenomenon and set of challenges." *Report of the American Psychological Association Task Force on the Interface Between Psychology and Global Climate Change*, [210]

UNEP. 2019. *Emissions Gap Report 2019*. United Nations Environment Programme. [206]

Valkengoed, Anne M. van, and Linda Steg. 2019. "Meta-analyses of factors motivating climate change adaptation behaviour." *Nature Climate Change* 9 (2): 158–63. [209]

Van Der Linden, Sander, Edward Maibach, John Cook, Anthony Leiserowitz, Michael Ranney, Stephan Lewandowsky, Joseph Árvai, and Elke U. Weber. 2017. "Culture versus cognition is a false dilemma." *Nature Climate Change* 7 (7): 457. DOI: 10.1038/nclimate3323. [220]

Welsch, Heinz. 2020. "Moral Foundations and Voluntary Public Good Provision: The Case of Climate Change." *Ecological Economics* 175: 106696. DOI: 10.1016/j.ecolecon.2020.106696. [210]

Xu, Yiqing, Jens Hainmueller, Jonathan Mummolo, and Licheng Liu. 2017. "INTERFLEX: Stata module to estimate multiplicative interaction models with diagnostics and visualization." *Statistical Software Components, Boston College Department of Economics*, (3). [224, 248]

Young, H. Peyton. 2008. "The New Palgrave Dictionary of Economics." In *The New Palgrave Dictionary of Economics*. Edited by Steven N. Durlauf and Lawrence E. Blume. London: Palgrave MacMillan. DOI: 10.1057/978-1-349-95121-5. [209]

Young, H. Peyton. 2015. "The Evolution of Social Norms." *Annual Review of Economics* 7 (1): 359–87. [209]

Appendix 3.A Supplementary analyses

Table 3.A.1. Comparison of the sample to the US population

Variable	Wave 1	Wave 2	ACS (2019)
Female	51%	51%	51%
Age: 18-34	30%	30%	30%
Age: 35-54	32%	32%	32%
Age: 55+	38%	38%	38%
Education: Bachelor's degree or above	32%	31%	31%
Region: Northeast	17%	17%	17%
Region: Midwest	21%	21%	21%
Region: South	38%	38%	38%
Region: West	24%	24%	24%

Notes: Columns 1 and 2 display the summary statistics for the survey samples of waves 1 and 2, respectively. Column 3 displays summary statistics based on the American Community Survey 2019.

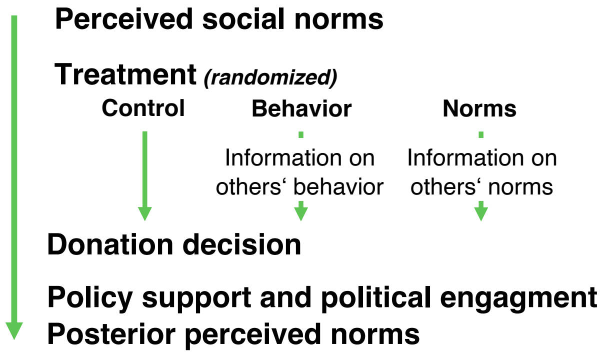

Notes: This figure provides an overview of the structure of the experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Figure 3.A.1. Structure of experiment

Table 3.A.2. Education and individual willingness to fight global warming

	Outcome: Donation ($)	
	(1)	(2)
Democrat x college degree	-6.838	-6.480
	(8.096)	(8.062)
Republican x college degree	-28.214***	-27.201***
	(10.320)	(10.429)
N	1,975	1,975
Control group mean	225.21	225.21
Demographic controls	Yes	Yes
Preferences and moral universalism	Yes	Yes
Normative belief	Behavior belief	Norms belief

Notes: This table shows OLS regression estimates where the dependent variable are donations (in $) using respondents from wave 1. All regressions specifications are identical to those in Table 3.3.1, including demographic controls, economic preferences, moral universalism as well as normative beliefs as covariates. However, we replaced the “College degree” indicator with a “Democrat x college degree” and a “Republican x college degree” indicator. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3.A.3. Test of balance

	Means (std. dev.)			Differences (p-values)		
	Control group (C)	Behavior treatment (T_1)	Norms treatment (T_2)	T_1 - C	T_2 - C	T_2 - T_1
Behavior belief	52.096	51.627	51.644	-0.470	-0.452	-0.017
	(21.339)	(21.213)	(21.391)	(0.486)	(0.503)	(0.980)
Norms belief	62.172	61.667	61.328	-0.505	-0.845	0.339
	(21.357)	(21.535)	(21.948)	(0.458)	(0.217)	(0.621)
Altruism	-0.008	-0.024	0.032	-0.016	0.040	-0.057*
	(0.982)	(0.984)	(1.032)	(0.600)	(0.206)	(0.076)
Patience	-0.020	0.005	0.015	0.025	0.035	-0.010
	(0.993)	(0.989)	(1.019)	(0.424)	(0.265)	(0.744)
Risk	-0.001	-0.005	0.006	-0.005	0.007	-0.011
	(0.989)	(1.011)	(1.000)	(0.887)	(0.827)	(0.719)
Pos. reciprocity	-0.018	0.021	-0.002	0.039	0.016	0.023
	(1.024)	(0.983)	(0.993)	(0.223)	(0.619)	(0.463)
Neg. reciprocity	-0.011	0.012	-0.001	0.023	0.011	0.013
	(0.999)	(0.978)	(1.023)	(0.455)	(0.733)	(0.692)
Trust	-0.028	0.017	0.010	0.045	0.038	0.007
	(1.001)	(1.000)	(0.999)	(0.156)	(0.229)	(0.825)
Rel. universalism	-0.027	0.021	0.006	0.047	0.032	0.015
	(0.987)	(1.020)	(0.993)	(0.138)	(0.303)	(0.639)
Age	48.114	47.350	47.847	-0.763	-0.266	-0.497
	(17.727)	(17.055)	(17.438)	(0.166)	(0.632)	(0.361)
Female	0.494	0.522	0.514	0.029*	0.020	0.008
	(0.500)	(0.500)	(0.500)	(0.071)	(0.202)	(0.593)
Log income	10.782	10.795	10.815	0.013	0.033	-0.020
	(0.882)	(0.879)	(0.858)	(0.645)	(0.236)	(0.471)
College degree	0.473	0.479	0.457	0.007	-0.015	0.022
	(0.499)	(0.500)	(0.498)	(0.676)	(0.335)	(0.166)
Employed	0.499	0.488	0.506	-0.012	0.007	-0.018
	(0.500)	(0.500)	(0.500)	(0.467)	(0.672)	(0.248)
Democrat	0.528	0.535	0.539	0.007	0.011	-0.003
	(0.499)	(0.499)	(0.499)	(0.640)	(0.497)	(0.833)
Northeast	0.170	0.165	0.174	-0.005	0.004	-0.009
	(0.376)	(0.372)	(0.380)	(0.692)	(0.717)	(0.447)
Midwest	0.204	0.211	0.216	0.007	0.012	-0.005
	(0.403)	(0.408)	(0.411)	(0.602)	(0.362)	(0.697)
South	0.390	0.385	0.365	-0.005	-0.025	0.020
	(0.488)	(0.487)	(0.482)	(0.743)	(0.105)	(0.196)
Parent	0.562	0.557	0.550	-0.005	-0.012	0.007
	(0.496)	(0.497)	(0.498)	(0.762)	(0.441)	(0.640)
p-value of joint F-test				0.426	0.684	0.425
Observations	1,987	1,995	2,018	3,982	4,005	4,013

Notes: Columns 1–3 show the means and standard deviations of respondent covariates in the different treatments of wave 2. Columns 4–6 show differences in means between the groups indicated in the column header together with p-values in parentheses. The p-values of the joint F-test are determined by regressing the treatment indicator on the vector of demographic controls. The F-test tests the joint hypothesis that none of the covariates predicts treatment assignment. *Covariates* "Behavior belief" and "Norms belief" are the perceived social norm measures, ranging from 0 to 100. Economic preferences (altruism, patience, risk, pos. reciprocity, neg. reciprocity, trust) and moral universalism (rel. universalism) are standardized. "Female", "Employed", "Democrat", "Parent", and the three census region dummies are binary indicators. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

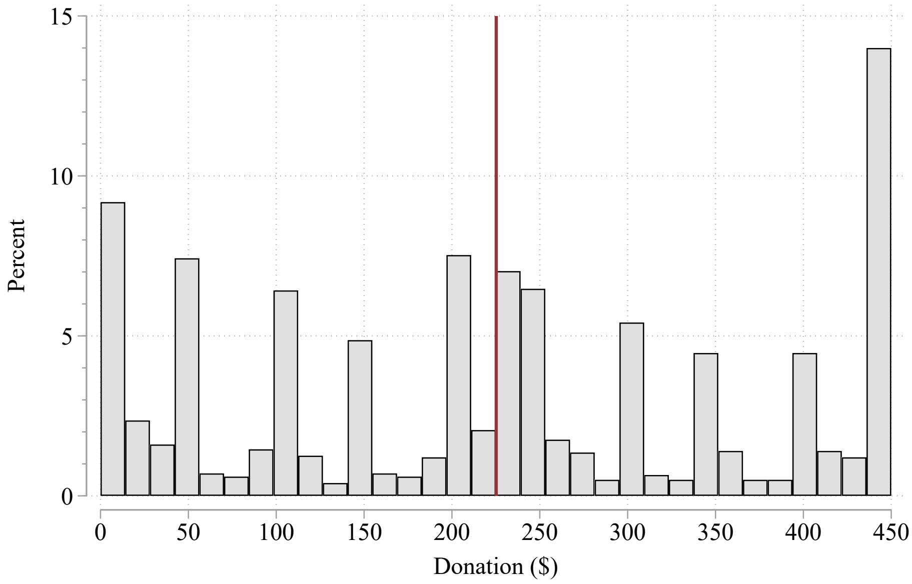

Notes: This figure shows the distribution of the monetary amounts donated to the climate charity in wave 1. The average donation is indicated by the vertical red line.

Figure 3.A.2. The distribution of individual willingness to fight global warming

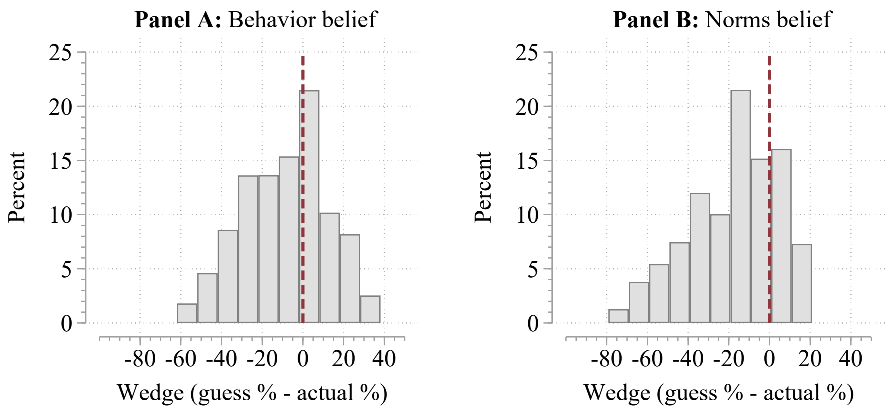

Notes: Using respondents from wave 2, this figure shows the distribution of the wedge between the respondent's perceived social norms and the actual shares in wave 1. Panel A shows people's belief about the share of Americans who say that they try to fight global warming. Panel B shows people's belief about the share of Americans who say that one should fight global warming. The red vertical line indicates the actual shares from wave 1.

Figure 3.A.3. Wedge in beliefs about social norms

Table 3.A.4. Determinants of norm misperceptions

	Dependent variable: Absolute prediction error (in percentage points)			
	Behavior belief		Norms belief	
	(1) Full sample	(2) Underestimators only	(3) Full sample	(4) Underestimators only
Democrat	-1.869***	-1.997**	-3.130***	-3.343***
	(0.663)	(0.868)	(0.814)	(0.945)
Age	0.063***	0.077***	0.133***	0.138***
	(0.021)	(0.027)	(0.026)	(0.030)
Female	0.919	0.823	1.378*	1.277
	(0.665)	(0.866)	(0.805)	(0.930)
Log household income	-0.508	-0.556	-1.104**	-1.617***
	(0.423)	(0.531)	(0.540)	(0.624)
College degree or more	-0.956	-0.264	-2.299***	-2.947***
	(0.727)	(0.969)	(0.892)	(1.050)
Currently employed	1.024	0.781	0.601	1.014
	(0.727)	(0.947)	(0.903)	(1.054)
Parent	-0.046	-1.238	-0.828	-0.745
	(0.703)	(0.915)	(0.863)	(0.998)
Constant	23.107***	26.513***	30.344***	39.914***
	(4.581)	(5.684)	(5.786)	(6.683)
N	1,996	1,334	1,996	1,519
R^2	0.013	0.013	0.033	0.040

Notes: This table shows OLS regression estimates using respondents from wave 1. The dependent variable in each column is the absolute difference between the respondent's stated belief (behavior/norms) and the actual share. "Behavior belief" is the respondent's belief about the share of Americans who fight global warming. "Norms belief" is the respondent's belief about the share of Americans who think one should fight global warming. Columns 1 and 3 use the full sample, while columns 2 and 4 focus on the subset of respondents who underestimate the actual shares. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

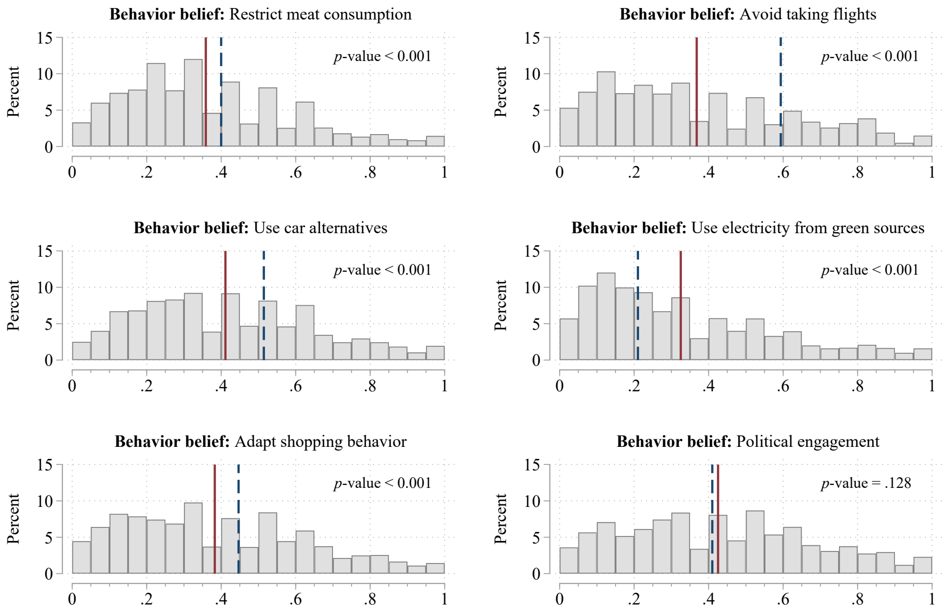

Notes: This figure shows the distribution of behavior beliefs in wave 1 for concrete climate-friendly behaviors. Each panel shows the distribution of people's beliefs about the share of Americans who say that they engage in the specific climate-friendly behavior indicated in the title of the panel. The solid red line indicates the average belief. The dashed blue line indicates the actual share of Americans engaging the behavior.

Figure 3.A.4. Perceived prevalence of concrete climate-friendly behaviors

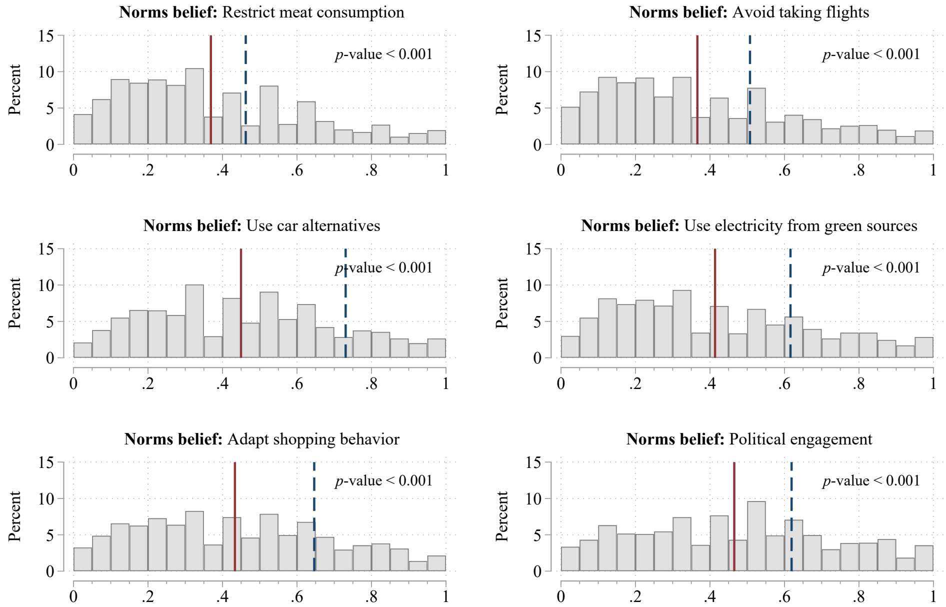

Notes: This figure shows the distribution of norms beliefs in wave 1 for concrete climate-friendly behaviors. Each panel shows the distribution of people's beliefs about the share of Americans who say that one *should* engage in the specific climate-friendly behavior. The solid red line indicates the average belief. The dashed blue line indicates the actual share of Americans saying that one should engage in the behavior indicated in the title of the panel.

Figure 3.A.5. Perceived prevalence of norms for concrete climate-friendly behavior

Table 3.A.5. Relationship of abstract and specific perceived norm measures

	(1) Restrict meat consumption	(2) Avoid taking flights	(3) Use car alternatives	(4) Use green electricity	(5) Adapt shopping behavior	(6) Political engagement
Panel A: Behavior						
Behavior belief	0.477***	0.362***	0.471***	0.421***	0.480***	0.468***
	(0.021)	(0.022)	(0.021)	(0.021)	(0.020)	(0.020)
N	1,994	1,993	1,993	1,994	1,992	1,993
R^2	0.228	0.131	0.222	0.178	0.231	0.219
Panel B: Norms						
Norms belief	0.410***	0.340***	0.454***	0.416***	0.471***	0.448***
	(0.021)	(0.021)	(0.021)	(0.020)	(0.020)	(0.020)
N	1,994	1,993	1,993	1,994	1,992	1,993
R^2	0.168	0.116	0.206	0.174	0.222	0.201

Notes: This table shows OLS regression estimates using respondents from wave 1. All coefficients can be interpreted as Pearson correlation coefficients. The dependent variables in Panel A are beliefs about the share of Americans who engage in the concrete climate-friendly behavior indicated in the column header. The dependent variables in Panel B are beliefs about the share of Americans who say that one should engage in the concrete climate-friendly behaviors. "Behavior belief" is the respondent's belief about the share of Americans who fight global warming. "Norms belief" is the respondent's belief about the share of Americans who think one should fight global warming. All beliefs are standardized to have a mean of zero and a standard deviation of one. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3.A.6. Treatment effects on climate donations and posterior beliefs: No controls

	(1) Donation ($)	(2) Behavior belief (post.)	(3) Norms belief (post.)
Behavior treatment	12.852***	0.285***	0.244***
	(4.824)	(0.031)	(0.031)
Norms treatment	17.485***	0.374***	0.355***
	(4.857)	(0.031)	(0.031)
N	5,991	5,988	5,976
Control group mean	249.31	0	0
z-scored	No	Yes	Yes

Notes: This table shows OLS regression estimates using respondents from wave 2. "Behavior treatment" is a binary indicator taking value one for respondents who received information about the share of Americans who try to fight global warming. "Norms treatment" is a binary indicator taking value one for respondents who received information about the share of Americans who say that one should try to fight global warming. "Behavior belief" is an index of six post-treatment beliefs about the share of Americans engaging in concrete climate-friendly behaviors to fight global warming. "Norms belief" is an index of six post-treatment beliefs about the share of Americans who say that one should engage in concrete climate-friendly behaviors to fight global warming. Both indices are standardized to have a mean of zero and a standard deviation of one in the control group. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3.A.7. Treatment effect heterogeneity: Climate change "denier"

	Dependent variable: Donation ($)				
	Interactant:				
	(1) No trust in science	(2) No scientific consensus	(3) Not concerned	(4) Not a threat	(5) Caused by nature
Panel A: Behavior treatment					
Treatment (a)	2.733	1.335	1.004	1.895	0.122
	(5.661)	(5.392)	(5.007)	(5.085)	(5.082)
Treatment x Interactant (b)	18.268*	22.561**	33.200***	29.943***	38.333***
	(9.357)	(10.126)	(10.410)	(10.330)	(10.466)
Interactant	-91.364***	-82.718***	-140.489***	-128.326***	-127.592***
	(7.145)	(7.472)	(7.751)	(7.710)	(7.865)
Linear combination (a + b)	21.001***	23.896***	34.204***	31.837***	38.455***
	(7.444)	(8.568)	(9.121)	(8.981)	(9.144)
N	3,978	3,978	3,978	3,978	3,978
Controls	Yes	Yes	Yes	Yes	Yes
Panel B: Norms treatment					
Treatment (a)	13.000**	8.245	10.241**	9.397*	11.639**
	(5.667)	(5.460)	(4.987)	(5.069)	(5.053)
Treatment x Interactant (b)	7.751	21.274**	14.928	14.560	14.569
	(9.353)	(10.044)	(10.406)	(10.398)	(10.386)
Interactant	-89.976***	-80.385***	-139.925***	-127.516***	-128.427***
	(7.140)	(7.465)	(7.742)	(7.726)	(7.852)
Linear combination (a + b)	20.751***	29.519***	25.169***	23.957***	26.208***
	(7.442)	(8.431)	(9.136)	(9.084)	(9.082)
N	4,000	4,000	4,000	4,000	4,000
Controls	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression from wave 2. The dependent variable is the donation to the climate charity ($). It is regressed on a treatment dummy for the behavior treatment (Panel A) and the norm treatment (Panel B), respectively, an interactant that varies across columns, and its interaction with the treatment dummy. Interactants are indicated by the column header. Each interactant is a binary variable taking value one. "No trust in science" means that the respondent trusts climate scientists "a moderate amount", "a little" or not at all (on a five-point Likert scale). "No scientific consensus" means that the respondent thinks that most scientists think that global warming is not happening or that there is no consensus among scientists. "Not concerned" means that the respondent is "not very worried' or "not at all worried" about global warming (on a four-point Likert scale). "Not a threat" means that the respondent thinks that global warming will do "only a little" or no harm at all (on a four-point Likert scale). "Caused by nature" means that the respondent thinks that global warming is caused by natural activities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3.A.8. Treatment effect heterogeneity: Climate change "denier" – Robustness to controlling for the interaction between treatment and prior beliefs

	Dependent variable: Donation ($)				
	Interactant:				
	(1) No trust in science	(2) No scientific consensus	(3) Not concerned	(4) Not a threat	(5) Caused by nature
Panel A: Behavior treatment					
Treatment (a)	9.683	9.228	12.670	13.362	13.353
	(13.391)	(13.147)	(12.671)	(12.741)	(12.617)
Treatment x Interactant (b)	17.090*	21.511**	32.559***	29.391***	38.440***
	(9.420)	(10.142)	(10.464)	(10.379)	(10.424)
Interactant	-89.111***	-80.718***	-138.804***	-126.678***	-127.584***
	(7.202)	(7.502)	(7.782)	(7.742)	(7.823)
Linear combination (a + b)	26.772**	30.739**	45.229***	42.753***	51.793***
	(13.242)	(13.961)	(13.740)	(13.726)	(13.967)
N	3,978	3,978	3,978	3,978	3,978
Controls	Yes	Yes	Yes	Yes	Yes
Treatment x Prior	Yes	Yes	Yes	Yes	Yes
Panel B: Norms treatment					
Treatment (a)	27.580*	18.851	22.250	25.774*	26.725*
	(15.650)	(15.657)	(15.042)	(15.231)	(14.881)
Treatment x Interactant (b)	5.596	18.748*	13.119	13.138	13.001
	(9.481)	(10.173)	(10.688)	(10.649)	(10.508)
Interactant	-84.081***	-74.126***	-134.167***	-121.945***	-123.874***
	(7.214)	(7.569)	(7.935)	(7.871)	(7.927)
Linear combination (a + b)	33.176**	37.599**	35.370**	38.912***	39.726***
	(14.660)	(15.099)	(14.536)	(14.745)	(14.640)
N	4,000	4,000	4,000	4,000	4,000
Controls	Yes	Yes	Yes	Yes	Yes
Treatment x Prior	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS regression from wave 2. The dependent variable is the donation to the climate charity ($). It is regressed on a treatment dummy for the behavior treatment (Panel A) and the norm treatment (Panel B), respectively, an interactant that varies across columns, and its interaction with the treatment dummy. Interactants are indicated by the column header. Each interactant is a binary variable taking value one. "No trust in science" means that the respondent trusts climate scientists "a moderate amount", "a little" or not at all (on a five-point Likert scale). "No scientific consensus" means that the respondent thinks that most scientists think that global warming is not happening or that there is no consensus among scientists. "Not concerned" means that the respondent is "not very worried' or "not at all worried" about global warming (on a four-point Likert scale). "Not a threat" means that the respondent thinks that global warming will do "only a little" or no harm at all (on a four-point Likert scale). "Caused by nature" means that the respondent thinks that global warming is caused by natural activities. All regressions include the corresponding prior belief and the interaction between the treatment indicator and the prior belief. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3.A.9. Preferences and universal values explain the partisan gap

	Donation ($)		Policy support	
	(1)	(2)	(3)	(4)
Democrat	74.323***	46.084***	0.923***	0.709***
	(6.523)	(6.279)	(0.041)	(0.040)
N	1,993	1,976	1,993	1,979
R^2	0.086	0.275	0.221	0.337
Demographic controls	Yes	Yes	Yes	Yes
Preferences and moral universalism		Yes		Yes

Notes: This table shows OLS regression estimates using respondents from wave 1. "Democrat" is a binary indicator taking value one if respondents identify with the Democrat party. We include our standard set of demographic controls: gender (indicator), age (continuous), log income, college degree (indicator), employment (indicator), and census region (three indicators). The dependent variable in columns 1–2 are donations, whereas the dependent variable in columns 3–4 is our standardized index of support for policies to fight global warming. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

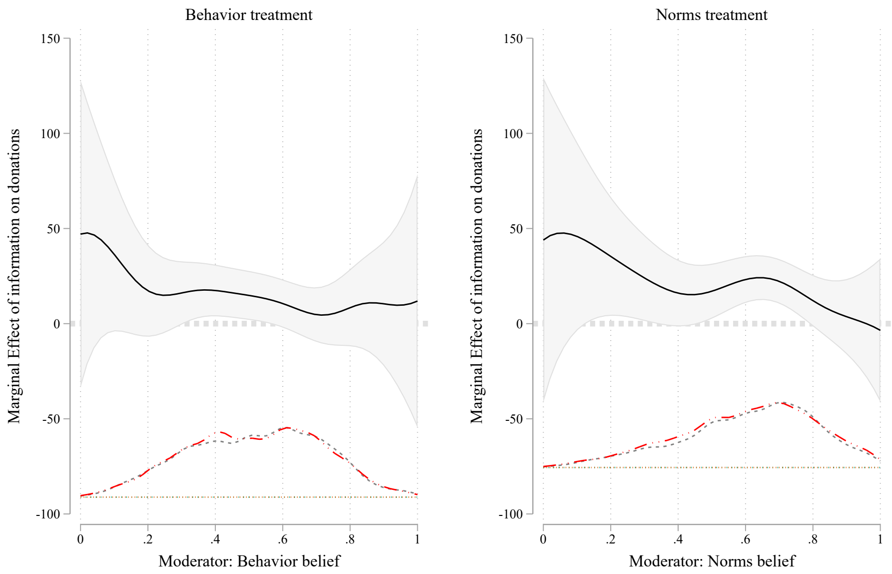

Notes: This figure shows the results from a non-linear interaction analysis using the *interflex* package (Xu et al., 2017; Hainmueller, Mummolo, and Xu, 2019) and restricting the sample to respondents from wave 2. The left panel excludes respondents in the norms treatment, while the right panel excludes respondents in the behavior treatment. The dashed lines at the bottom of each panel plot the distribution of the pre-treatment belief. 95% confidence intervals using robust standard errors are shown. Both panels show results without including additional controls.

Figure 3.A.6. Treatment effect heterogeneity by perceived social norms: Non-parametric estimates

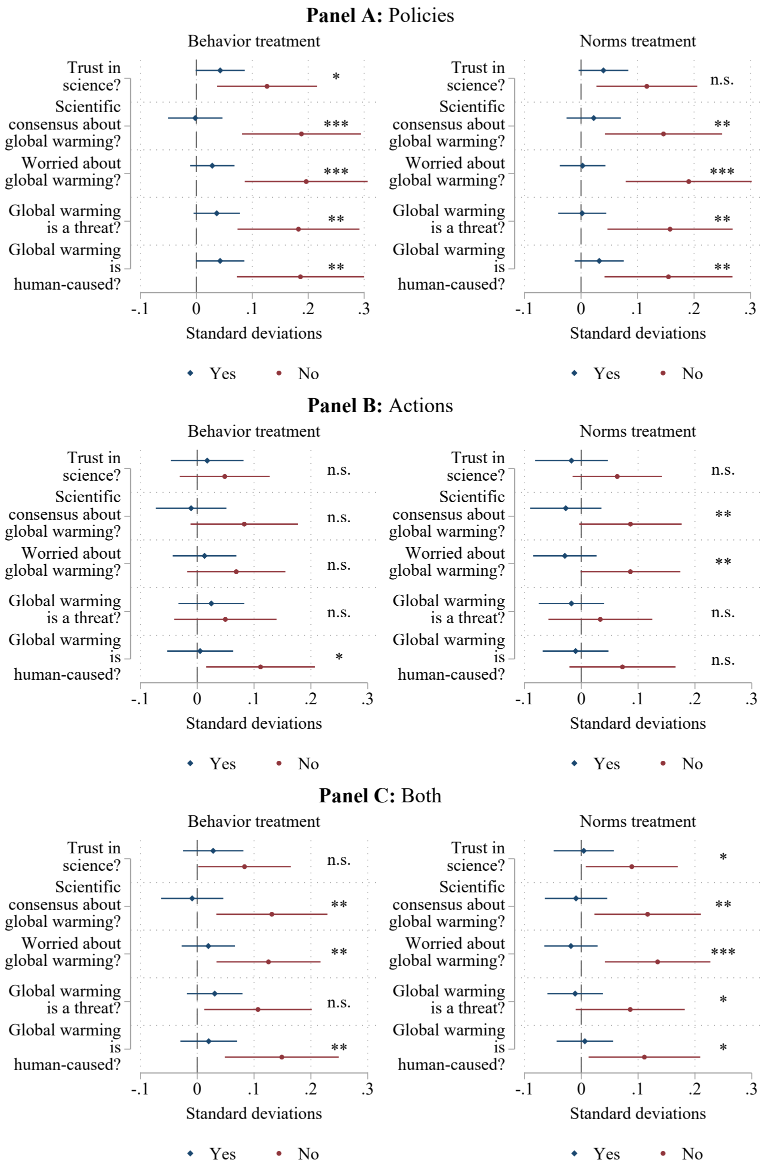

Notes: This figure shows treatment effects in different subsamples using respondents from wave 2. Panel A shows treatment effects on the policy support index, Panel B shows treatment effects of the action index, and Panel C shows treatment effects on the joint index. 95% confidence intervals are shown. Each panel shows estimates for the subsample of climate change deniers – e.g., those who have no trust in science or do not believe in human-caused global warming – and the subsample of respondents who are not skeptical of climate change. "Trust in science" means that the respondent trust climate scientists "a lot" or "a great deal" (on a five-point Likert scale). "Scientific consensus about global warming" means that the respondent thinks that most scientists think that global warming is happening. "Worried about global warming" means that the respondent is "somewhat worried' or "very worried" about global warming (on a four-point Likert scale). "Global warming is a threat" means that the respondent thinks that global warming will do "a moderate amount" or "a great deal" of harm (on a four-point Likert scale). "Global warming is human-caused" means that the respondent thinks that global warming is caused by human activities.

Figure 3.A.7. Heterogeneity by "climate change denier": Political outcomes

Appendix 3.B Questionnaire

This appendix presents the main survey blocks, following the order of exposition in the paper. The full questionnaire containing all questions administered as part of this study can be downloaded from https://osf.io/chvy6/.

3.B.1 Attention screener

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This compromises the results of research studies. **To show that you are reading the survey carefully, please choose both "Very strongly interested" and "Not at all interested" as your answer to the next question.** Given the above, how interested are you in politics?

1.Very strongly interested
2.Very interested
3.A little bit interested
4.Not very interested
5.Not at all interested

Only participants who select both (a) and (e) pass this attention screener.

3.B.2 Measuring individual willingness to fight climate change

A decision about money

Please pay special attention to the next question in which you will make a decision about money. We will randomly select 25 respondents. If you are among them, your decision will be a real decision. The decision will be implemented and you can receive up to $450.

Your decision

Here is the decision: You can divide $450 between yourself and a charitable organization that fights global warming. The amount that you keep for yourself will be added to your account. The amount that you donate will go to the award-winning charity *atmosfair*. *atmosfair* actively contributes to CO_2 mitigation by promoting, developing and financing renewable energies worldwide. In this way, a donation saves CO_2 that would otherwise be created by fossil fuels. *atmosfair* spends around $12 million per year to fight global warming and uses less than 5% of donated funds to cover administrative costs. You can find more information on *atmosfair* here.

It costs about \$450 to offset the yearly CO_2 emissions of a typical US citizen. This number is calculated as follows: It costs about \$28 to prevent 1 ton of CO_2 emissions. The World Bank estimates that a typical US citizen causes about 16 tons of CO_2 emissions per year.

How much of the \$450 would you like to donate to *atmosfair*?

3.B.3 Introducing bonus scheme

Bonus payment possible

There are several questions in this survey, in which we will ask you to guess how other respondents answered a question. These questions are flagged with the sign:

You can earn a bonus of \$1. This works as follows: We will randomly select one of the flagged questions. Your response to this question is considered as correct if it differs at most by three from the correct number you are asked to guess. If your response to this question is correct, \$1 will be added to your account.

3.B.4 Measuring perceived social norms

Do you try to fight global warming?
[Yes/No]
Do you think that people in the United States should try to fight global warming?
[Yes/No]

[PAGE BREAK]

The questions on this page are bonus questions. This means that you can earn additional money if you answer them correctly.

As part of this research project, we recently surveyed many people in the United States and asked them the same questions. Respondents come from all parts of the population and their responses represent the views and attitudes of people in the United States.
What do you think? Out of 100 people we asked, how many stated that...

1.... they try to fight global warming?

2.... they think that people in the United States should try to fight global warming?

3.B.5 Treatments: Shifting perceived social norms

3.B.5.1 Behavior treatment

What do other people in the United States do?

We recently surveyed 2,000 people in the United States and asked them whether they try to fight global warming. Respondents come from all parts of the population and their responses represent the views and attitudes of people in the United States. On the next page, you will learn how they responded. Please read the information carefully.

[PAGE BREAK]

We asked 2,000 Americans: Do you try to fight global warming? Yes or no?
Here are the results:

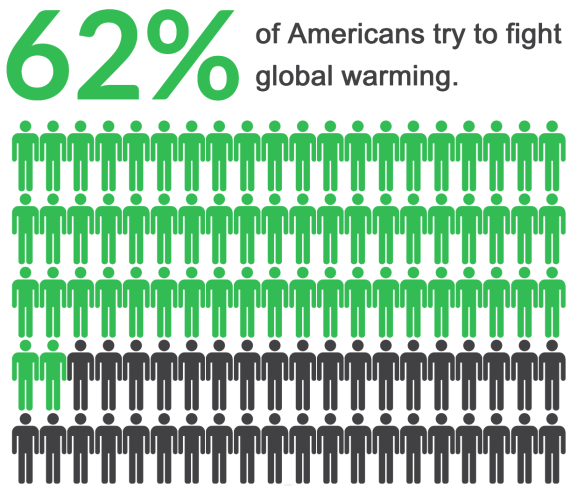

3.B.5.2 Norms treatment

What do other people in the United States think?

We recently surveyed 2,000 people in the United States and asked them whether they think people in the US should try to fight global warming. Respondents come from all parts of the population and their responses represent the views and attitudes of people in the United States. On the next page, you will learn how they responded. Please read the information carefully.

[PAGE BREAK]

We asked 2,000 Americans: Do you think that people in the United States should try to fight global warming? Yes or no?
Here are the results:

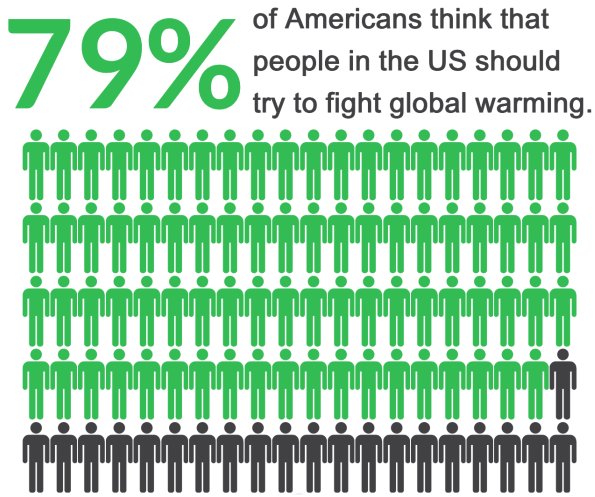

3.B.6 Measuring posterior beliefs

The questions on this page are bonus questions. This means that you can earn additional money if you answer them correctly.

BONUS

As part of this research project, we recently surveyed many people in the United States and asked them the same questions. Respondents come from all parts of the population and their responses represent the views and attitudes of people in the United States.
We asked respondents to state whether they have taken different actions to fight global warming over the last year.
What do you think? Out of 100 people we asked, how many stated that...

1.... restrict their meat consumption?

2.... avoid taking flights?

3.... regularly use environmentally-friendly alternatives to their private car such as walking, cycling, taking public transport or car-sharing?

4.... receive electricity only from green/renewable sources (e.g., solar energy or wind power)?

5.... adapt their shopping behavior to the carbon footprint of products?

6.... politically support the fight against global warming, e.g. participate in a demonstration, sign a letter, or support a political organization?

[PAGE BREAK]

Do you think that people in the United states **should**...

1.... restrict their meat consumption?

2.... avoid taking flights?

3.... regularly use environmentally-friendly alternatives to their private car such as walking, cycling, taking public transport or car-sharing?

4.... receive electricity only from green/renewable sources (e.g., solar energy or wind power)?

5.... adapt their shopping behavior to the carbon footprint of products?

6.... politically support the fight against global warming, e.g. participate in a demonstration, sign a letter, or support a political organization?

[PAGE BREAK]

The questions on this page are bonus questions. This means that you can earn additional money if you answer them correctly.

BONUS

What do you think? Out of 100 people we asked the same questions, how many stated that they think that people in the United States should...

1.... restrict their meat consumption?

2.... avoid taking flights?

3.... regularly use environmentally-friendly alternatives to their private car such as walking, cycling, taking public transport or car-sharing?

4.... receive electricity only from green/renewable sources (e.g., solar energy or wind power)?

5.... adapt their shopping behavior to the carbon footprint of products?

6.... politically support the fight against global warming, e.g. participate in a demonstration, sign a letter, or support a political organization?

3.B.7 Measuring climate change skepticism

In general, how much do you trust scientists who do research on global warming?

1.A great deal

2.A lot

3.A moderate amount

4.A little

5.Not at all

Which comes closest to your own view?

1.Most scientists think global warming is happening.

2.There is a lot of disagreement among scientists about whether or not global warming is happening.

3.Most scientists think global warming is not happening.

How worried are you about global warming?

1.Very worried

2.Somewhat worried

3.Not very worried

4.Not at all worried

How much do you think global warming will harm people in the United States?

1.Not at all

2.Only a little

3.A moderate amount

4.A great deal

Do you think that global warming is mainly...?

1.a result of human activities

2.a result of natural causes

3.B.8 Measuring policy support and political engagement

Taken from the detailed politics module developed as part of the Climate Change in the American Mind Project (Howe et al., 2015).

Policy support

How much do you support or oppose the following policies?
Strongly support / Somewhat support / Somewhat oppose / Strongly oppose

1.Fund more research into renewable energy sources, such as solar and wind power.

2.Regulate carbon dioxide (the primary greenhouse gas) as a pollutant.

3.Set strict carbon dioxide emission limits on existing coal-fired power plants to reduce global warming and improve public health. Power plants would have to reduce their emissions and/or invest in renewable energy and energy efficiency. The cost of electricity to consumers and companies would likely increase.

4.Require fossil fuel companies to pay a carbon tax and use the money to reduce other taxes (such as income tax) by an equal amount.

5.Require electric utilities to produce at least 20% of their electricity from wind, solar, or other renewable energy sources, even if it costs the average household an extra $100 a year.

6.Provide tax rebates for people who purchase energy-efficient vehicles or solar panels.

How much do you agree or disagree with the following statements?
Strongly agree / Somewhat agree / Somewhat disagree / Strongly disagree

1.Schools should teach our children about the causes, consequences, and potential solutions to global warming.

Political engagement

How likely would you be to do each of the following things?
Definitely would / Probably would / Probably would not / Definitely would not

1.Vote for a candidate for public office because of their position on global warming.

2.Publicly display t-shirt, bumper sticker, button, wrist band, or sign about global warming.

3.Donate money to an organization working on global warming.

4.Volunteer your time to an organization working on global warming.

5.Write letters, email, or phone government officials about global warming.

6.Meet with an elected official or their staff about global warming.

7.Support an organization engaging in non-violent civil disobedience against corporate or government activities that make global warming worse.

8.Personally engage in non-violent civil disobedience (e.g., sit-ins, blockades, or trespassing) against corporate or government activities that make global warming worse.

9.Attend a political rally, speech, or organized protest about global warming.

10.Write a letter to the editor of a newspaper or magazine or call a live radio or TV show to express an opinion about global warming.

11.Share information about global warming on social media.

Appendix 3.C Construction of variables

3.C.1 Measuring economic preferences

We administer the Global Preferences Survey (GPS) and follow the methodology described in Falk, Becker, Dohmen, Enke, et al. (2018) to obtain detailed individual-level measures of economic preferences. More information on the construction of the variables can be found below.

1. *Patience.* The measure of patience (or time preference) is derived from the combination of responses to two survey measures, one with a quantitative and one with a qualitative format. The quantitative survey measure consists of a series of five interdependent hypothetical binary choices between immediate and delayed financial rewards. In each of the five questions, participants have to decide between receiving a payment today or a larger payment in 12 months. The qualitative measure of patience is given by the respondents' self-assessment regarding their willingness to wait on an eleven-point Likert scale, asking "how willing are you to give up something that is beneficial for you today in order to benefit more from that in the future?".
2. *Risk Taking.* Risk preferences are also elicited through a series of related quantitative questions as well as one qualitative question. Just as with patience, the quantitative measure consists of a series of five binary choices. Choices are between a fixed lottery, in which the individual could win x or zero, and varying sure payments, y. The qualitative item asks for the respondents' self-assessment of their willingness to take risks on an eleven-point scale ("In general, how willing are you to take risks?").
3. *Positive Reciprocity.* Positive reciprocity is measured using one quantitative item and one qualitative question. First, respondents are presented a choice scenario in which they are asked to imagine that they got lost in an unfamiliar area and that a stranger – when asked for directions – offered to take them to their destination. Respondents are then asked which out of six presents (worth between 10 and 60 dollars) they would give to the stranger as a "thank you". Second, respondents are asked to provide a self-assessment about how willing they are to return a favor on an eleven-point Likert scale.
4. *Negative Reciprocity.* Negative reciprocity is elicited through three self-assessments. First, respondents are asked how willing they are to take revenge if they are treated very unjustly, even if doing so comes at a cost (Likert scale, 0-10). The second and third items probe respondents about their willingness to punish someone for unfair behavior, either towards themselves or a third person.
5. *Altruism.* Altruism is measured through a combination of one qualitative and one quantitative item, both of which are related to donations. The qualitative question asks respondents how willing they would be to give to good causes

without expecting anything in return on an eleven-point scale. The quantitative scenario depicts a situation in which the respondent unexpectedly receives 1,600 dollars and is asked to state how much of this amount they would donate.

6. *Trust.* The trust measure is based on one item, which asks respondents whether they assume that other people only have the best intentions (Likert scale, 0-10).

For each economic preference, the survey items are combined into a single preference measure. More specifically, each preference is computed by (i) calculating the z-scores of each survey item at the individual level and (ii) weighting these z-scores using the weights provided in Table 3.C.1. For ease of interpretation, each preference measure is standardized to have a mean of zero and a standard deviation of one.

Table 3.C.1. GPS Survey Items and Weights

Preference	Item description	Weight
Patience	Intertemporal choice sequence using staircase method	0.712
	Self-assessment: willingness to wait	0.288
Risk taking	Lottery choice sequence using stair case method	0.473
	Self-assessment: willingness to take risks in general	0.527
Positive reciprocity	Gift in exchange for help	0.515
	Self-assessment: willingness to return a favor	0.485
Negative reciprocity	Self-assessment: willingness to take revenge	0.374
	Self-assessment: willingness to punish unfair behavior toward self	0.313
	Self-assessment: willingness to punish unfair behavior toward others	0.313
Altruism	Donation decision	0.635
	Self-assessment: willingness to give to good causes	0.365
Trust	Self-assessment: people have only the best intentions	1

3.C.2 Measuring universal moral values

Moral Foundation Theory posits that people's moral concerns can be split into five foundations:

1. *Care/Harm*. This foundation measures the extent to which people care about the weak and try to keep others away from harm.
2. *Fairness/Reciprocity*. This measure captures the importance of equality, justice, rights and autonomy.
3. *In-group/Loyalty*. This foundation captures the extent to which people emphasize loyalty to the "in-group" (family, country) and how morally relevant betrayal is.
4. *Authority/Respect*. This foundation measures how important respect for authority, tradition and order is.
5. *Purity/Sanctity*. This measure captures the importance of ideas related to purity, disgust and traditional religious attitudes.

To obtain measures of the five foundations, we administer the Moral Foundations Questionnaire. In this survey, each moral foundation is measured using six different survey items. Respondents are either asked to assess the moral relevance of certain behaviors, or they are asked if they agree with certain moral value statements. All the questions are answered on a Likert scale (0–5). Table 3.C.2 provides an overview of the specific items that are included in each foundation. In order to construct the final scores, responses are summed.

To construct a measure of the relative importance of universal versus communal moral values, we follow the approach described in Enke (2020):

$$\begin{aligned} &\text{Relative importance of universal values} && (3.C.1)\\ &= \text{Universal values} - \text{Communal values} && (3.C.2)\\ &= \text{Harm/Care} + \text{Fairness/Reciprocity} - \text{In-group/Loyalty} - \text{Authority/Respect} && (3.C.3) \end{aligned}$$

To ease interpretation, the resulting measure is standardized to have a mean of zero and a standard deviation of one.

Table 3.C.2. Survey items: Moral Foundations Questionnaire

	Moral Relevance	Agreement with Statement
Harm/care	Emotional suffering	Compassion with suffering crucial virtue
	Care for weak and vulnerable	Hurt defenseless animal is the worst thing
	Cruelty	Never right to kill human being
Fairness/reciprocity	Treat people differently	Laws should treat everyone fairly
	Act unfairly	Justice most important requirement for society
	Deny rights	Morally wrong that rich children inherit a lot
In-group/loyalty	Show love for country	Proud of country's history
	Betray group	Be loyal to family even if done something wrong
	Lack of loyalty	Be team player, rather than express oneself
Authority/respect	Lack of respect for authority	Children need to learn respect for authority
	Conform to societal traditions	Men and women have different roles in society
	Cause disorder	Soldiers must obey even if disagree with order
Purity/sancity	Violate standards of purity	Not do things that are disgusting
	Do something disgusting	Call acts wrong if unnatural
	Act in a way that God would approve	Chastity is an important virtue

Note: For the items in column 1, respondents are asked to state to what extent these considerations are morally relevant (Likert scale from 0 to 5). For the items in column 2, respondents are asked to state whether they agree or disagree with the statements (Likert scale from 0 to 5).

Chapter 4

Intertemporal Altruism

Joint with Philipp Eisenhauer, Armin Falk, and Thomas Graeber

Abstract: Most prosocial decisions involve intertemporal tradeoffs. Yet, the *timing* of prosocial utility flows is ambiguous and has largely disregarded in models of other-regarding preferences. We study the behavioral implications of the time structure of prosocial utility, leveraging a conceptual distinction between *consequence-dated* and *choice-dated* utility flows. We conduct a high-stakes donation experiment that comprehensively characterizes discounting behavior in self-other tradeoffs and allows us to identify different prosocial motives from their distinct time profiles. Our data can only be explained by a combination of choice- and consequence-dated prosocial utility. Both motives are pervasive and negatively correlated at the individual level.

Acknowledgements: The experiments in this paper were approved by the Ethics Committee of the Economics Department at the University of Bonn (reference no. 2016-02). For helpful comments and discussions we are grateful to Peter Andre, Alexander Cappelen, Bertil Tungodden, Christine Exley, Florian Zimmermann, James Konow, Johannes Abeler, Jörgen Weibull, George Loewenstein and Nicola Lacetera. We thank Annica Gehlen and Emily Schwab for their outstanding research assistance. Armin Falk acknowledges financial support from the German Research Foundation through the Leibniz Program as well as European Research Council Project 340950—Institutions and Morality.

4.1 Introduction

In prosocial decisions, choices and consequences are typically separated in time. Donations, for example, tend to create immediate costs to the donor and delayed benefits for others. Consider climate change charities, which routinely face the choice between promoting either adaptation or mitigation projects. The benefits from adaptation projects tend to accrue much earlier than those from mitigation projects. If individuals only care about the timing of the donation itself, then the different planning horizons of such projects should not affect their willingness to contribute. If, on the other hand, individuals do care about the timing of benefits, then charities are well-advised to take the different time frames into account. Consider instead a commitment to voluntary work where both the costs to the donor and the benefits to others are delayed. Similarly, repeated interactions such as reciprocal exchange also naturally involve intertemporal considerations. I may expect to reciprocate a favor from someone else later on, trading off an earlier benefit against a delayed cost. The inherent intertemporal nature of prosocial choices raises important questions about how choice environments affect the timing and level of prosocial choices and how we should think about the timing of the utility flows associated with prosocial decisions.

Notably, the existing theoretical literature on prosocial preferences largely abstracts from the time dimension of utility flows. For example, outcome-based models of inequity aversion do not specify how to evaluate inequality that occurs across two points in time (see, e.g. Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). To illustrate, consider a simple donation or dictator game with a delayed payment to the recipient. Do inequity-averse donors discount the corresponding recipient's utility in the same way as they discount their utility? Do their social preferences apply to the discounted utility stream (of self and recipient), or do they care about period-specific inequality? These timing-related considerations are not unique to inequity aversion, but apply to other forms of social preferences alike. In formal models of reciprocity (Charness and Rabin, 2002; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006), social interactions are conceptualized as being inherently timeless. Returning a favor one year later is considered just as worthwhile as returning a favor now. From a different perspective, the concept of warm glow (Andreoni, 1989, 1990) explicitly suggests that utility may derive from the act of choice itself rather than the prosocial externality. However, the corresponding theories do not distinguish between the timing of choice and delayed consequences. Similarly, models of image concern (e.g., Bénabou and Tirole, 2006) do not specify whether image utility accrues at the point of prosocial choice or at the time of its consequences or observability. The common practice of modeling prosocial behavior as *atemporal* limits our scope for understanding prosocial behavior in practice, which typically features a separation of choices and consequences over time as in the above examples. This gap in the literature calls for more discipline on the role of delays in theoretical and empirical work on prosocial behavior.

We provide a theoretically guided empirical investigation of discounting behavior in a high-stakes donation context. Unlike related empirical work, we do not focus on partial delays in dictator games (Kovarik, 2009; Dreber, Fudenberg, Levine, and Rand, 2016), the role of commitment (Rogers and Bazerman, 2008; Breman, 2011), or time inconsistency and present bias (Kölle and Wenner, 2018; Andreoni and Serra-Garcia, 2021). Instead, our experimental approach allows us to characterize entire discount functions in self-other trade-offs in a highly comprehensive and novel manner. Our analysis proceeds in three steps. First, we develop a conceptual distinction between *consequence-dated* and *choice-dated* utility in modeling intertemporal prosocial choice. This formal distinction leverages existing theoretical and empirical work in a productive fashion and provides a fruitful guiding framework for our own empirical exercise. If utility is consequence-dated, then it accrues with a delay that corresponds to when the actual utility consequences for others materialize. If utility is choice-dated, then it is realized in temporal proximity to the act of giving. We derive qualitative predictions of models with choice- and consequence-dated utility in different contexts. Second, we conduct a controlled laboratory study and establish a set of reduced-form patterns in atemporal and intertemporal donation behavior that directly speak to our model predictions. Third, we implement a structural model and estimate an explicit intertemporal utility function that reproduces the core patterns in our data and allows us to assess the relative importance of consequence-dated and choice-dated utility in determining prosocial behavior. Our experiment is deliberately designed to provide transparent identification of the different utility components (in the spirit of, e.g., DellaVigna, 2018).

To experimentally study the intertemporal dimension of prosocial choice in a meaningful way, we implement a choice paradigm with far-ranging real-world implications. In our incentivized, high-stakes donation paradigm, each participant could save human lives by individually causing donations of up to 800 euros for the treatment of tuberculosis patients by a designated charity and earn up to 200 euros for themselves. The unusually high incentives serve to make both the donation context and the implemented delays meaningful to subjects. For all choice tasks, we use a variant of the widely used multiple price list methodology. The experiment comprises two parts: a series of intertemporal choice tasks in which participants decide between dated certain payments to themselves or the charity for delays of up to twelve months, and a series of atemporal risky choice tasks to characterize participants' multi-attribute utility function representing preferences over "self-euros" and "charity-euros." The first part is further divided into three stages. Across stages, we vary whether choices present (a) trade-offs between earlier and later payments in a single utility domain (only self-euros or only charity-euros), (b) trade-offs between payments in different domains that involve a unique, common payment date either now or in the future, and (c) trade-offs across domains and payment dates that require self-other comparisons across time. This setup systematically examines behavior when either (a) only time matters, (b) only cross-attribute comparisons matter, or (c) both time and cross-attribute comparisons matter. To our knowledge, this is the first experiment providing data that are rich enough to allow for sharp tests of the discounted utility model in the multi-attribute case of self-other tradeoffs.

We purposefully opted for a design with monetary pay-offs because (i) *prosocial* utility flows are not typically associated with primary consumption by the decision-maker such as food or effort; (ii) we aim to characterize discount functions comprehensively, including for time horizons in excess of one month, which has not been accomplished with real-effort designs so far;[1] and (iii) our interest is partly in the application to monetary donations, which is the most widespread form of altruistic behavior in practice and has immediate consequences for charities. The recent methodological review by Cohen, Ericson, Laibson, and White (2019) discusses situations in which money designs may be preferable to real-effort paradigms, which we argue includes our case of studying prosocial utility flows that are typically not yoked to primary consumption by the decision-maker.[2] Our design deliberately abstracts from the issue of present bias and the phenomenon of extreme short-run impatience by implementing payments as wire transfers. Even the soonest possible experimental payment was subject to a delay of three days, which the literature conventionally considers as being "in the future."[3]

We start with a discussion of our reduced-form findings and document non-parametric evidence compatible with consequence-dated as well as choice-dated prosocial utility. First, in smaller-sooner, larger-later choices involving either only self-euros or only charity-euros, subjects discount both delayed self-euro and delayed charity-euro payments. The notion that delayed donations are less valuable to subjects implies that valuations of charity-euros are linked to their payment date, pointing towards the existence of a consequence-dated component of prosocial utility flows. This qualitative devaluation pattern of delayed donations obtains for *all*

1. Real effort experiments have been conducted for short-time horizons of up to a few weeks for logistical reasons that mainly concern trust issues and attrition (Augenblick and Rabin, 2018; Augenblick, 2019).

2. Outside of the topic of other-regarding preferences, it has been pointed out that money designs may confound the timing of payments with the timing of primary consumption (Cubitt and Read, 2007; Chabris, Laibson, Morris, Schuldt, and Taubinsky, 2008). The emergent view in this literature may be that subjects tend to treat money like consumption (perhaps due to narrow bracketing), except in very short-time horizons (Halevy, 2015; Augenblick, 2019; Cohen et al., 2019; Balakrishnan, Haushofer, and Jakiela, 2020).

3. There are alternative methodologies, including the recent innovation of convex time budgets (Andreoni and Sprenger, 2012). While convex time budgets do not require a separate estimation of the utility function, we prefer the "double multiple price list" method of characterizing the atemporal utility function using separate choices (e.g. Andersen, Harrison, Lau, and Rutström, 2008). In so doing, we can examine the features of the multi-attribute atemporal utility function in more detail and circumvent the issue of bunching at the boundaries and choice inconsistencies frequently observed with convex time budgets (Chakraborty, Calford, Fenig, and Halevy, 2017).

intertemporal decisions that involve a time trade-off between the two choice options, including cross-attribute intertemporal decisions. More strikingly, net present values measured for delayed self-euros and delayed charity-euros are statistically indistinguishable. Non-parametric analyses imply that our combined data from atemporal choices and choices involving time trade-offs are specifically in line with the discounted utility specification of consequence-dated utility, i.e., an intertemporal utility function that applies the same discount function to future utility streams generated by self-euro and charity-euro payments. Second, however, when contemporaneous, identically-dated self-euro and charity-euro payments are delayed into the future, subjects become increasingly more willing to give up self-euros for charity-euros as the delay increases. These choices that create a cross-attribute but no time trade-off imply a declining subjective exchange rate between charity-euros and self-euros. To our knowledge, we provide the first dataset that allows documenting such a pattern based on experimental variation. This finding is incompatible with a stationary flow utility function as posited by the discounted utility model where identically-dated utility flows are subject to the same discount factor. Under those circumstances, the effect of discounting cancels out, and we expect a constant exchange rate. Instead, our finding of a declining forward exchange rate suggests that the prosocial utility derived from donating money has a choice-dated component that is not subject to discounting due to, for example, warm glow or self-image concerns. We can only rationalize a declining subjective exchange rate if prosocial utility from donating (partly) accrues at the time of choice and is independent of the timing of the actual payment. Hence, our reduced-form findings suggest both a consequence-dated and a choice-dated component of prosocial utility. However, none of the existing models of prosocial behavior are compatible with this combination of motives.

We fill this gap and develop a simple model of intertemporal prosocial choice that accommodates both consequence-dated and choice-dated prosocial utility flows. We fit this model to our data using structural estimations at both the population and subject level. Our structural analysis adds two insights. First, our estimated structural model replicates the distinctive choice patterns identified in our reduced-form analysis. Most importantly, we are able to replicate a declining forward exchange rate because the relative weight of choice-dated utility in the discounted prosocial utility increases. As choice-dated utility is not discounted, the overall prosocial utility thus declines less quickly in the delay than the discounted utility from equally-delayed self-euros. Our parameter estimates for standard preferences parameters are in line with previous work. Second, the structural analysis sheds light on the individual-level variation of parameters, revealing that the different forms of prosociality display marked heterogeneity. We find that 80% of subjects exhibit meaningful, positive consequence-dated prosociality, and just below 60% of subjects show meaningful, positive choice-dated prosociality. Strikingly, there is a strong negative correlation between the two parameters at the subject level. This negative relationship indicates that differently-dated prosocial motivations might characterize distinct "types" of subjects. Some are primarily driven by consequence-dated motives such as pure altruism, whereas others seem to follow choice-date motivations such as image concerns or the feeling of warm glow.

We build on and contribute to several stands of the literature. Our conceptual distinction between consequence-dated and choice-dated prosocial motives complements existing research on what motivates contributions to public goods and charitable giving. While departing from existing work in terms of our focus on the time dimension rather than – for example – the impact of one's generosity and the corresponding "neutrality" hypothesis (Andreoni, 1989), we view the distinction drawn here as a natural extension and re-interpretation of the work on warm glow and pure altruism. Focusing exclusively on intertemporal arguments leads us to conclude the existence of mixtures of both motives, which resonates with previous work that documents mixed motivations, i.e. "impure altruism" (Andreoni, 1993; Bolton and Katok, 1998; Konow, 2010).[4] The distinction between choice-dated and consequence-dated prosocial utility provides a productive framework to extend models of prosocial behavior to an intertemporal context. It predicts that the primary motivation for prosocial behavior *changes* with the temporal delay. While considerations of consequence will be more important when they are realized in temporal proximity, the choice-dated component of prosocial utility will drive choices involving consequences that are strongly separated in time. This *switch* implies that simply extrapolating previous evidence on the relative importance of different prosocial motives from atemporal contexts to intertemporal settings may lead to inaccurate conclusions.

We also provide the first comprehensive experimental dataset on intertemporal prosocial behavior using a fully-crossed design of choices involving single vs. cross-attribute trade-offs – self-euro vs. charity-euro payments – and short vs. long delays. The concept of a "forward exchange rate" characterizes behavior for increasing, *common* delays, which provides a non-parametric test of the discounted utility model. Accordingly, our experimental approach allows us to address questions about the nature of intertemporal prosocial trade-offs that cannot be answered with a subset of this data. Previous empirical research has focused on different aspects of intertemporal self-other tradeoffs as outlined above (Rogers and Bazerman, 2008; Kovarik, 2009; Breman, 2011; Dreber et al., 2016; Kölle and Wenner, 2018; Andreoni and Serra-Garcia, 2021). While our account rationalizes some of this evidence through the implied time patterns of flow utility rather than, e.g., a hyperbolic shape of the discount function, we view our work as fruitfully complementing this emerging body

4. Moreover, our finding of correlation aversion – i.e., that the marginal utilities of self-euro and charity-euro payments are not independent - leads to the substantive interpretation that own earnings and donations are partial substitutes. This underscores the emerging consensus on a relationship between income, wealth, and charitable giving (Meer and Priday, 2020).

of evidence that has different objectives and focuses on different phenomena such as time inconsistency and present bias.

Additionally, our findings inform work on intertemporal multi-attribute utility more generally. The literature has only recently started to explore the ramifications of multi-attribute utility functions for modeling intertemporal choice (Andersen, Harrison, Lau, and Rutström, 2018). Although related empirical work studies the patterns of multi-attribute, intertemporal choices (Cubitt, McDonald, and Read, 2018), it only looks at typical consumption goods rather than self-other trade-offs and – unlike our paper – does not quantify the effects using structural estimation. While our results from single-domain discounting choices are in line with a unique, domain-general discount function, which is a key assumption of the discounted utility model, previous studies report discounting patterns that sometimes differ across goods (Chapman, 1996; Frederick, 2006; McClure, Ericson, Laibson, Loewenstein, and Cohen, 2007; Hardisty and Weber, 2009; Kim, Schnall, and White, 2013). These studies have different objectives from ours, and consequently, they do not separately account for the shape of the atemporal utility function and do not rely on high-stakes experimental designs.

The paper proceeds as follows. Section 4.2 lays out a theoretical framework for our argument. Section 4.3 describes the experimental design and procedures. We present our reduced-form results in Section 4.4 and the structural analyses in Section 4.5. Finally, Section 4.6 concludes.

4.2 Conceptual framework

We develop a simple formal framework that is not intended as an exhaustive theoretical characterization of intertemporal prosocial choice and that we do not consider a main contribution of the paper. Instead, the objective of this section is twofold: first, it carves out our main conceptual distinction between consequence- and choice-dated utility flows in a tractable and easily generalizable fashion. Second, the framework disciplines and guides our subsequent empirical analysis.

Standard theory assumes that individuals derive utility from the consumption of goods and services. However, prosocial choices such as donations are usually not associated with primary rewards and require additional assumptions about the sources of utility. Consequently, research in psychology allows for a broader notion of consumption that is not limited to physical consumption but instead involves forms of conceptual consumption that occur entirely in the mind (Schelling, 1988; Ariely and Norton, 2009). In line with this approach, the economic literature on prosocial preferences puts forward a variety of motives such as intentions (Falk and Fischbacher, 2006) or image concerns (Bénabou and Tirole, 2006) that are independent of immediate consumption by the decision-maker. This variety of prosocial motivations naturally lends itself to distinguish between the time structures of corresponding utility flows. We apply the canonical notion of dated period utility from intertemporal choice theory but disentangle two constituent elements of prosocial behavior. We introduce an explicit distinction between the act of making a prosocial choice and the consequences of this choice for others. In this framework, we refer to utility flows as *choice-dated* if they are realized at the time of giving and as *consequence-dated* if they accrue when the consequences for others actually materialize.

We seek to understand what this conceptual distinction implies for intertemporal prosocial choice and what we can learn from observed choices about the nature of prosocial preferences. In a first step, we address these questions and discuss the implications of models in which decision-makers receive *only* consequence-dated prosocial utility or *only* choice-dated prosocial utility. Our setting is deliberately simplified as we explore the two extreme cases to explore and contrast their distinct implications. In a second step, we consider the mixed case where both types of prosocial utility are present.

Let t index the current period in which a choice is made, and τ denote the time relative to the choice period. Let $x_{t+\tau}$ represent a dated payment to the decision-maker at time $t+\tau$ ("self-euros"). Moreover, let $g_{t+\tau}$ denote a dated payment to a charity at time $t+\tau$ ("charity-euros"). The decision-maker has preferences over dated payment streams $z = (x_{t+\tau}, g_{t+\tau})_{\tau\in\mathbb{N}}$ represented by an intertemporal utility function $U(z)$. We do not assume a specific form of prosocial preferences at this stage and treat self-euros $x_{t+\tau}$ and charity-euros $g_{t+\tau}$ as direct inputs to the utility function. In line with the previous literature (Halevy, 2015; Balakrishnan, Haushofer, and Jakiela, 2020), payment dates serve as a proxy for the conversion of money into utility for the self or others or as (sufficiently delayed) monetary payments treated as consumption goods. To simplify the following analysis, we interpret payment dates as representing the corresponding consumption dates.

We specify payments to others as a direct input into the utility function of the decision-maker. This approach is consistent with the interpretation that the decision-maker's prosocial utility truly depends on the utility – rather than the pay-off – consequences for others. Our conclusions remain unchanged as long as the recipient's utility is monotonic in the payments that they receive and approximated in time by the payment dates. Thus, we refrain from specifying the recipient's utility function for simplicity.[5]

In the spirit of providing conceptual guidance for our experimental study, we now separately discuss the concepts and psychological motivations of consequence-dated and choice-dated utility separately before contrasting their empirical predictions.

5. If we assume that the other person's utility is – ceteris paribus – a monotone function $v(g)$ of donations, we can substitute $v(g)$ for g in the utility function and study the reduced form.

4.2.1 Consequence-dated prosocial utility

In the case of consequence-dated prosocial utility, the utility of a donation to charity $g_{t+\tau}$ at time $t+\tau$ will also accrue at $t+\tau$, even if caused by a choice at an earlier point in time t. In this case, choosing between two dated payments to a charity with different payment dates requires an intertemporal comparison of prosocial utility flows. We can draw on standard economic tools and assume that the decision-maker behaves as if she maximizes her discounted intertemporal utility. The following intertemporal utility function then characterizes models of consequence-dated prosocial utility:

$$U_t = \sum_{\tau=0}^{T} D(\tau)u(x_{t+\tau}, g_{t+\tau}). \tag{4.2.1}$$

We make the standard assumptions that there is a stationary discount function $D(\tau)$ that applies to future utility flows (Cohen et al., 2019). The flows are represented by a stationary flow utility function, $u(x_{t+\tau}, g_{t+\tau})$, which captures the decision-maker's concern for herself and others.

Two remarks about this specification are in order. First, while we remain deliberately neutral about the precise psychological motives underlying consequence-dated prosocial utility, pure altruism provides a natural interpretation of Equation (4.2.1). A pure altruist cares about the welfare consequences of their choices, which in the model is determined by $g_{t+\tau}$. Any self-other trade-off then involves interpersonal utility comparisons, suggesting the interpretation of u as the decision-maker's subjective welfare function for evaluating contemporaneous consequences of her choices to the self and others. Second, a complementary perspective on the intertemporal utility function in Equation (4.2.1) is the natural extension of the workhorse model of intertemporal choice – discounted utility – to the multi-attribute case, because it conceptualizes self-euros and charity-euros as conventional arguments of a flow utility function. Consequently, the interpretation of prosocial behavior in an intertemporal context through the lens of multi-attribute discounted utility is akin to adopting the perspective of consequence-dated prosocial utility.

4.2.2 Choice-dated prosocial utility

In the case of choice-dated prosocial utility, the utility of a dated donation $g_{t+\tau}$ accrues in the period t in which it was caused through a choice, even if the payment is executed at a later date $t+\tau$. This implies that earlier and later donations to charity generate the same utility to the decision-maker. It introduces a theoretical distinction between consequence-dated and choice-dated prosocial utility that allows us to obtain sharp qualitative descriptions.[6] We can then represent choice-dated prosocial

6. It is possible that delayed donations provide lower choice-dated prosocial utility. However, choice-dated prosocial utility should devalue at a *lower rate* than consequence-dated utility, as it is

utility with the following intertemporal utility function:

$$U_t = \sum_{\tau=0}^{T} \alpha\left(g_{t+\tau}\right) + \sum_{\tau=0}^{T} D(\tau) u(x_{t+\tau}), \tag{4.2.2}$$

where $\alpha\left(g_{t+\tau}\right)$ is the choice-dated and immediate prosocial utility that results from causing a potentially delayed donation today. Note that for our illustrative purposes here, we rule out complementarities between self-euros and charity-euros as well as interactions between choice-dated utility derived from actions with consequences that materialize with different delays.[7] Again, we do not take a stance on the psychological motives of choice-dated utility and its specific relationship to the size of a donation. However, our formulation naturally encompasses a wide range of motives. They include the feeling of warm glow that is explicitly defined as being related to the act of giving (Andreoni, 1989, 1990) and self- or social-image concerns that are routinely characterized as being linked in time to the act of donating rather than to the instrumental value of charitable funds.

4.2.3 Qualitative predictions

We contrast the implications of models of choice-dated and consequence-dated prosocial utility for intertemporal choices involving self-euros and charity-euros.

In Figure 4.2.1, each axis represents one of the following three trade-offs: (1) pure time trade-offs (univariate discounting, UD_τ), (2) pure across-domain trade-offs (subjective exchange rates, F_τ) and (3) mixed across-time and across-domain trade-offs (multivariate discounting, MD_τ).[8]

---

otherwise indistinguishable from consequence-dated considerations. This means that as the delay increases, the prosocial motivation in choice-dated models will be relatively more stable compared to the prosocial motivation in consequence-dated models. Our results only require this *relative* property. To simplify the exposition, we directly assume that choice-dated utility is independent of the delay.

7. One could accommodate these complementarities using more general classes of utility functions such as $U_t = V(W(a) + b)$ where a and b represents the two sums in Equation (4.2.2).

8. In Appendix 4.C, we discuss the case of choice-dated prosocial utility more extensively under weaker assumptions and obtain qualitatively similar predictions.

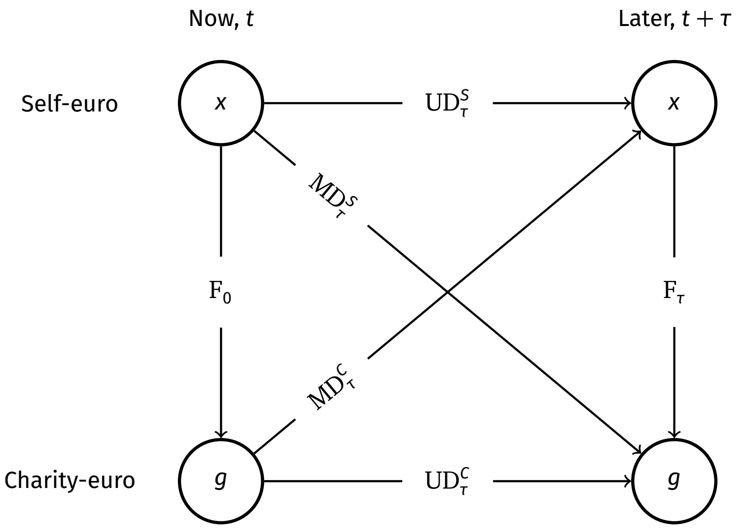

Figure 4.2.1. Intertemporal self-other trade-offs

Notes: This figure displays three intertemporal self-other trade-offs.

We begin with the horizontal axes in Figure 4.2.1, which capture the standard case of univariate discounting (UD_τ). A decision-maker can choose between receiving m_t charity-euros (self-euros) at time t or receiving a larger payment of $m_{t+\tau}$ charity-euros (self-euros) at a later time $t+\tau$. The prediction of consequence-dated prosocial utility is that the value of charity-euros (self-euros) decreases by $D(\tau)$ with the additional delay τ. While choice-dated prosocial utility necessarily makes the same qualitative prediction for univariate discounting of self-euros, the immediate gratification from giving other-euros to charity is not subject to discounting.

Prediction 3. Delayed charity-euros are discounted in consequence-dated models, but not in choice-dated models, of prosocial behavior. Both models predict discounting of delayed self-euros.

Next, we turn to the vertical axes in Figure 4.2.1 and consider the exchange rate F_τ, which describes the decision-maker's subjective conversion rate between contemporaneous self-euros and charity-euros in τ periods. It is defined as $F_\tau = g^*_{t+\tau}/x_{t+\tau}$ whenever the decision-maker is indifferent between $g^*_{t+\tau}$ and $x_{t+\tau}$.[9] In a choice-dated model, the corresponding indifference condition is

$$D(\tau)u(x_{t+\tau},0) \;=\; D(\tau)u(0,g^*_{t+\tau}). \tag{4.2.3}$$

9. The exchange rate will depend on the level of payments unless the utility function satisfies homogeneity, but we omit the dependence for ease of exposition.

As the discount factor $D(\tau)$ cancels from this expression, the exchange rate F_τ does not depend on τ. Note that this holds irrespective of the shape of the flow utility function, providing the distinctive prediction of a constant exchange rate for models of consequence-dated prosocial utility. By contrast, in models of choice-dated prosocial utility, the defining equation of the exchange rate takes the following form:

$$D(\tau)u(x_{t+\tau}) \;=\; \alpha(g^*_{t+\tau}). \tag{4.2.4}$$

As the delay τ of both payments increases, the decision-maker discounts the value of self-euros on the left-hand side, while the choice-dated prosocial utility remains unaffected. Thus, $g^*_{t+\tau}$ decreases, causing the exchange rate F_τ to decrease in τ.

Prediction 4. Consequence-dated models predict a constant exchange rate, whereas choice-dated models of prosocial behavior predict a declining exchange rate.

Finally, we turn to the diagonal axes in Figure 4.2.1, which capture multivariate discounting (MD_τ). Similar to the exchange rate, this intertemporal trade-off only arises in the multi-attribute case. A decision-maker receives m_t self-euros (charity-euros) at time t and is then asked to state the dated payment $m_{t+\tau}$ of charity-euros (self-euros) to be received at a later time $t+\tau$ that makes her indifferent. This decision involves a choice between payments to different recipients at different points in time, and provides an implicit multivariate discount factor of $m_t/m_{t+\tau}$. As in the case of univariate discounting, consequence-dated models will discount the value of the later payment, irrespective of whether it is denominated in self-euros or charity-euros. In both cases, we expect to see multivariate discounting. If the earlier payment involves self-euros, the indifference condition is $u(m_t,0)=D(\tau)u(0,m_{t+\tau})$. The right-hand side decreases with τ, while the left-hand side is constant, causing multivariate discounting. In the other case, we have the symmetric condition $u(0,m_t)=D(\tau)u(m_{t+\tau},0)$. For models of choice-dated prosocial utility, we obtain the same prediction of multivariate discounting only when the early payment is denoted in charity-euros, because then the value of delayed self-euros is also discounted. However, we expect no multivariate discounting if the early payment involves self-euros. The reason is again that the immediate, choice-dated prosocial utility is unaffected by the delay τ of charity-euros. The indifference condition is as follows:

$$u(m_t,0) \;=\; \alpha(m_{t+\tau}).$$

Prediction 5. Consequence-dated models predict multivariate discounting, whereas choice-dated models of prosocial behavior predict multivariate discounting if the later payment involves self-euros and no multivariate discounting if the later payment involves charity-euros.

Figure 4.2.1 summarizes the predictions that we now explore in our tailored experimental setting. It is straightforward to obtain qualitative predictions for the mixed case of both choice-dated and consequence-dated prosocial utility.

Table 4.2.1. Predictions of different models

	Type of prosocial utility		
Prediction	Choice-dated	Consequence-dated	Both
Univariate discounting of self- and charity-euros		✓	✓
Declining exchange rate	✓		✓
Multivariate discounting for both self- and charity-euros as today's numeraire		✓	✓

4.3 Experimental design and procedures

We set up a tightly controlled experiment that allows the precise manipulation of payment dates, including a credible implementation of future payments and donations. At the same time, the stakes remain quantitatively meaningful even when payments are delayed substantially.

4.3.1 Saving a Life donation paradigm

To make delays in experimental outcomes relevant to subjects, our design attempts to take prosocial decision-making in a controlled setting to the limits: we developed a high-stakes donation paradigm in cooperation with the Indian non-profit organization Operation ASHA, which specializes in the treatment of tuberculosis, the world's deadliest bacterial infectious disease (World Health Organization, 2020).[10] Operation ASHA's model for treating tuberculosis has received extensive public acclaim and worldwide media coverage. Under conservative assumptions, a donation of 350 euros – roughly 400 US dollars at the time – covered all costs incurred by Operation ASHA to identify, treat and cure five patients, which is equivalent to saving one additional human life in expectation.[11]

10. See Operation ASHA's website at http://www.opasha.org for details.

11. We estimated the all-inclusive cost of a life saved by Operation ASHA based on public information on the charity's operations in combination with estimates from peer-reviewed epidemiological studies on tuberculosis mortality (Kolappan, Subramani, Kumaraswami, Santha, and Narayanan, 2008; Straetemans, Glaziou, Bierrenbach, Sismanidis, and Werf, 2011; Tiemersma, Werf, Borgdorff, Williams, and Nagelkerke, 2011). We conferred our donations as a restricted grant ensuring that no money is used to cover overhead costs and that the donations flow immediately into scaling up the Operation ASHA's treatment model.

Our experimental instructions provided detailed information about the causes, prevalence, and implications of tuberculosis and Operation ASHA.[12] All information on tuberculosis was verifiable and came from acknowledged sources, in particular the *World Health Organization*. We directly transferred all donations to Operation ASHA's bank account on the exact day specified in the experiment and offered subjects the opportunity to inspect proof of the bank transfer.

4.3.2 Design

The experiment comprises two consecutive parts: intertemporal choices (Part A) and atemporal choices under risk (Part B). Across both parts of the experiment, each subject completed a total of 36 decision screens, 21 involving intertemporal choices and 15 involving choices under risk. In each part, one randomly-chosen row of the price list on a randomly-chosen decision screen was selected by the computer and added to the subject's earnings. Before we provide the implementation details on both parts, two general remarks about the experimental design are in order.

First, we implement choices involving monetary payments to the subjects and the charity, rather than primary consumption such as effort or food. While most research on discounting behavior has relied on financial rewards, the recent experimental literature emphasizes that the discounted utility model posits discounting of *utility*, and that monetary payments only enter utility via primary consumption. Cohen et al. (2019) review this literature and conclude that studies using financial flows tend to find lower discount rates and a less hyperbolic discount function, implying smaller present bias. In the present study on self-other trade-offs in the context of donations, we use monetary payments, because most donations in practice are denominated in money. Our interest lies in time horizons exceeding two months, which has previously not been studied using primary consumption due to the logistical complications. Furthermore, we aim to circumvent the issue of genuine present bias to identify choice-related utility flows. The differences between discounting of financial flows and primary consumption are most pronounced for very early rewards, and previous work has argued that monetary rewards that do not occur in the immediate future are treated as consumption (Halevy, 2015; Augenblick, 2019; Balakrishnan, Haushofer, and Jakiela, 2020). Building on this debate, our deliberate design choice of avoiding utility consequences from consumption "in the present" allows for the simplifying assumption that delayed payments directly enter the utility function. In our setting, even the earliest payment date in our experiment lies "in the future". Specifically, we execute payments as bank transfers, with the earliest payment being available to subjects no sooner than three days following the day of the experiment.

12. See our experimental instructions in the Appendix.

Second, we use the widely-established multiple price list method for all intertemporal and risky choice tasks (Schubert, Brown, Gysler, and Brachinger, 1999; Holt and Laury, 2002; Attema, Bleichrodt, Gao, Huang, and Wakker, 2016; Dohmen, Falk, Huffman, and Sunde, 2017). On each decision screen, subjects faced a list of binary decisions between a fixed left-hand-side amount and a right-hand-side option with increasing amounts from the top to the bottom of the list. It is well established in the intertemporal choice literature that estimates of discount rates from simple "money earlier versus later" choices alone are confounded given pervasive evidence against linear utility even for small amounts. Several approaches address this issue (Montiel Olea and Strzalecki, 2014; Ericson and Noor, 2015), including the recently popular paradigm of convex time budgets, which does not require a separate elicitation of the utility curvature (Andreoni and Sprenger, 2012). We instead rely on the "double price list method", which estimates the shape of the atemporal utility function from separate risky choices, extending the approach of Andersen et al. (2008) to the multi-attribute case. While both methods have been shown to perform well in practice (Andreoni, Kuhn, and Sprenger, 2015), we primarily resort to using separate risky choices due to our objective of precisely characterizing the multi-attribute atemporal utility function.[13]

4.3.2.1 Part A – Intertemporal choices

We study intertemporal choices involving payments of self-euros and charity-euros by implementing a fully-crossed design with decisions involving cross-attribute vs. no cross-attribute trade-offs and differential delays vs. no differential delays. Using multiple price lists as shown in Appendix Figure 4.A.1, we elicit indifference points between certain self-euro or charity-euro payments at different, exactly-specified delays. Part A comprises five stages presented in randomized order.

Univariate discounting includes two stages, *UD – SELF* and *UD – CHARITY*, in which we separately elicit net present values of delayed payments of self-euros or charity-euros, respectively. On each decision screen of stage *UD – SELF*, subjects face a list of binary choices between a fixed payment of 50 self-euros to be received by bank transfer at the earliest possible payment date after three days and increasing amounts of self-euros at a fixed later point in time. The delay of the later payment varies across decision screens and may be either 1, 3, 6, or 12 months, in randomized

13. Note that both methods have practical disadvantages. While choices from convex time budgets produce substantial bunching at the boundaries and choice inconsistencies (Chakraborty et al., 2017), the price list methodology creates a substantial minority of subjects who switch multiple times in a single list, which is at odds with monotonic preferences (e.g. Bruner, 2011). Here we circumvent the complications associated with multiple switching points in the data by enforcing a unique switching point. This was implemented using an auto-completion function that filled in remaining choices as soon as a subject switched from the fixed left-hand-side option to the increasing right-hand-side option.

order. Subjects complete four decision screens in stage *UD – SELF*. Stage *UD – CHARITY* is identical to *UD – SELF* except that both the earlier and later payments involve donations to charity, which would be made by bank transfer on the specified dates in a way that could be verified by subjects later on. In our univariate discounting choices, individuals face a trade-off between two payments for the *same* recipient (either self-euros or charity-euros) that occur at *different* points in time.

We measure subjective exchange rates between self-euro and charity-euro payments at different points in time in stage *ER*. On each decision screen, subjects face a list of binary choices between a payment of 50 self-euros at a specified point in time and increasing amounts of charity-euros at the *same* point in time. Time points include bank transfers to be expected with the shortest delay of three days (the *spot exchange rate*) as well as in 1, 3, 6, or 12 months (*forward exchange rates*). These five decision screens provide measures of how many charity-euros subjects demand per contemporaneous self-euro for different delays from today's perspective. Note that the choices about the subjective exchange rate present individuals with trade-offs between two payments for *different* recipients, but occurring at the *same* points in time.

We measure trade-offs between two payments – one denominated in self-euros and one in charity-euros – with different delays. Stages *MD – SELF* and *MD – CHARITY* thus capture the common situation in which individuals face trade-offs between giving and taking, but the corresponding payment flows occur at different times. On each decision screen in stage *MD – SELF*, subjects face a list of binary choices between a fixed payment of 50 self-euros at the earliest delay and increasing amounts of charity-euros at a fixed later point in time. Conversely, in stage *MD – CHARITY*, subjects face a list of binary choices between a fixed payment of 50 charity-euros at the earliest delay and increasing amounts of self-euros at a fixed later point in time. As before, the later time points include 1, 3, 6 and 12 months. Multivariate discounting choices create trade-offs between two payments for *different* recipients, occurring at *different* points in time.

Within Part A, both the order in which stages occur and the order of decisions within each stage are randomized at the individual level.[14] Right-hand-side options in the price lists range from a simple annualized discount rate of 0% to 150% in increments of five percentage points for univariate discounting, from zero euros to 200 euros in increments of 10 euros for the exchange rates, and from zero euros to an annualized discount rate of 150% (relative to the 50 euros left-hand-side option) in 25 steps in stages *MD – CHARITY* and *MD – SELF*.

14. To avoid confusion, all decision screens belonging to the same stage appeared consecutively (in randomized order).

4.3.2.2 Part B – Risk apportionment

The objective of Part B is to characterize individuals' multi-attribute utility functions using atemporal decisions, i.e. choices that do not involve differently-dated payments. Note that the intertemporal choices in Part A only identify discounting behavior under the assumptions that flow utility is linear and additively separable in its attributes.

We adopt the recently popularized experimental paradigm of risk apportionment, which allows for non-parametric testing conditions on the nature of the utility function. Second- and third-order risk aversion (i.e. prudence) are typically defined in terms of specific conditions on the (second and third) derivatives of the utility function under expected utility maximization. Eeckhoudt and Schlesinger (2006) provide an alternative definition based on observable choices in risk apportionment tasks. Risk apportioning has the desirable feature that the measurement remains valid even if expected utility theory fails (Starmer, 2000; Ebert and Kuilen, 2015). At the same time, data from risk apportionment choices allow us to calibrate specific utility specifications under additional parametric assumptions.

We measured univariate risk aversion individually for self-euros and for charity-euros (stages *RA – SELF* and *RA – CHARITY*, respectively), univariate prudence (stages *PR – SELF* and *PR – CHARITY*), and multivariate risk aversion (stage *X – RA*). The latter stage is crucial as it delivers a non-parametric estimate of *correlation aversion* (Richard, 1975; Epstein and Tanny, 1980), which is a sufficient condition for assuming additive non-separability of the utility function.

In every risk apportionment task, subjects receive some endowment $\boldsymbol{e} = (x,y)$ of attributes X and Y and then make a decision between two lotteries. Each of these lotteries has two equally likely outcomes. Assume further that there are two undesirable fixed amounts R_1 and R_2 with $R_i \preceq (0,0)$. Accordingly, R_1 is a fixed univariate "reduction" in either X or Y, but not in both dimensions at the same time.[15] A preference for risk apportionment is the desire to disaggregate these unavoidable fixed reductions in wealth, R_1 and R_2, across two equiprobable states of the world, as depicted in Figure 4.3.1.

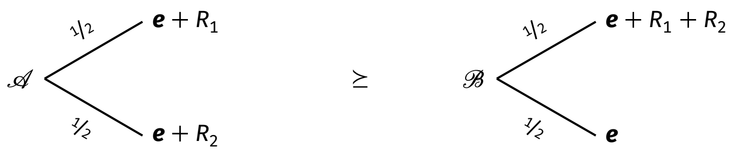

Figure 4.3.1. Preference for risk apportionment (cf. Ebert and Kuilen (2015))

The different stages in Part B vary depending on whether each attribute (X and Y) corresponds to self-euros or charity-euros. Concretely, we present subjects

15. The same holds for R_2, but R_1 and R_2 do not necessarily affect the same attribute.

with choices between two lotteries as summarized in Figure 4.3.1. For conceptual consistency and to avoid confusing subjects, we employ the same price list methodology as for intertemporal choices in Part A.[16] On each decision screen, subjects make binary choices between a fixed lottery $\mathscr{A}$ and a fixed lottery $\mathscr{B}$, where an additional, state-independent compensation payment m is added to lottery $\mathscr{B}$. This compensation payment m gradually increases across the rows of the choice list. The smallest amount for which the individual prefers lottery $\mathscr{B}$ indicates the minimal compensation demanded for heaving both undesirable reductions in wealth clustered in a single state. An example choice screen is depicted in Appendix Figure 4.A.2.

Table 4.3.1. Overview of risk apportionment choices

Stage	Endowment		R_1		R_2		Expected value	
	Self	Charity	Self	Charity	Self	Charity	Self	Charity
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
RA – SELF	25		-10		-5		17.5	
	50		-20		-10		35	
	100		-40		-20		70	
PR – SELF	40		-10		(14, 0.5; -14, 0.5)		35	
	40		-10		(7, 0.8; -28, 0.2)		35	
	40		-10		(-7, 0.8; 28, 0.2)		35	
RA – CHARITY		25		-10		-5		17.5
		50		-20		-10		35
		100		-40		-20		70
PR – CHARITY		40		-10		(14, 0.5; -14, 0.5)		35
		40		-10		(7, 0.8; -28, 0.2)		35
		40		-10		(-7, 0.8; 28, 0.2)		35
X – RA	25	25	-10			-10	20	20
	50	50	-20			-20	40	40
	100	100	-40			-40	80	80

Note: All values are displayed in euros. Columns labeled “Self” indicate payments to the subject and columns labeled “Charity” indicate payments to the charity. If R_1 or R_2 is a non-degenerate lottery, it is given as $(x_1, p_1;\ x_2, p_2)$, where x_i indicates the amount and p_i the probability of receiving it. Columns 8 and 9 show the expected payment to the subject and the expected payment to the charity, respectively.

Table 4.3.1 shows all fifteen choice scenarios presented to subjects. Note that for our measure of prudence, R_2 is a zero-mean lottery instead of a fixed reduction in wealth, i.e. R_2 only adds variance in this case. The grid of compensations offered in the choice lists varies with the endowments. Each choice list contains 21 rows across which the compensation increases at equal intervals. All grids are centered at zero.

16. Concretely, our design extends the procedure suggested in Ebert and Wiesen (2014) to a multi-attribute setting.

4.3.2.3 Procedures

We recruited 244 subjects from the student subject pool of the *BonnEconLab* at the University of Bonn. Table 4.A.1 provides summary statistics for the full sample. The experiment was conducted in the main auditorium at the University of Bonn. We collected data in nine sessions from September 19 to September 22, 2016. The experiment was fully computerized and conducted using the software oTree (Chen, Schonger, and Wickens, 2016). Subjects were seated in separate cubicles to create full privacy so that no other person could see their screen during the experiment. They could ask questions to an experimenter at any time. The average completion time was 65 minutes.

Subjects received a fixed amount of five euros for their participation in the experiment. All payments were made as bank transfers initiated on the precise day indicated for the payment. On average, each participant earned 59 euros (39 euros at the earliest delay and 20 euros at later time points) and caused donations of 70 euros (40 euros at the earliest delay and 30 euros at later time points). Average earnings and average donations together corresponded to fifteen times the federal hourly minimum wage at the time, or more than 10% of the median monthly household income in our sample.

4.3.3 Transformations

We ensure comparability of the compensation payments in each lottery and divide them by their expected value. To make intertemporal choices comparable across tasks, we proceed as follows. For choices from the stages UD and MD, we calculate the net present value (expressed in today's numeraire) of a dated future payment of one euro from subjects' smaller-sooner-larger-later choices. Specifically, the net present value is $50/m^*$, where m^* is the subject's switching point.[17] For choices from stage ER, we calculate the (forward) exchange rate $m^*/50$, i.e. the rate of charity-euros per contemporaneous self-euro.

4.4 Reduced-form results

We now document the main qualitative patterns in our data. Our goal is to disentangle the consequence-dated and choice-dated models of prosocial utility by testing the distinctive predictions of the two models. We first outline our estimation strategy and then discuss our results.

17. The value of the earlier payment (option A) in the multiple price list is always 50 euros. We use the midpoint of the interval where the subjects switched from option A to option B.

4.4.1 Estimation strategy

We present the average net present values and exchange rates by task. We use non-parametric hypothesis tests (Wilcoxon signed-rank test, paired t test) for inference about differences in means. These tests exploit the within nature of our design and ignore the between-subject variation in choices. The construction of our 95% confidence intervals is based on the procedure developed by Morey (2008) and Cousineau (2005). The procedure is best understood by considering the following auxiliary regression analysis of our results. Let $y_{i,j}$ denote an outcome of interest derived from subject i's selection of task j. We then estimate the saturated regression model separately for the stages UD, MD, and ER:

$$y_{i,j} = \alpha_i + \beta \text{Domain}_j + \sum_{\tau} \gamma_\tau \text{Delay}_{\tau(j)} + \sum_{\tau} \delta_\tau \text{Domain}_j \times \text{Delay}_{\tau(j)} + \varepsilon_{i,j}. \tag{4.4.1}$$

Here, α_i is a subject fixed effect, Domain_j is a binary variable taking the value of 1 if the earlier dated payment in task j is denoted in charity-euros, $\text{Delay}_{\tau(j)}$ is a binary variable taking the value of 1 if the later dated payment in task j has a delay of τ months, and $\varepsilon_{i,j}$ denotes the individual error term. To account for the nature of the within design, we cluster standard errors at the subject level.

The confidence intervals developed by Morey (2008) and Cousineau (2005) for differences in means across tasks will be similar to the confidence intervals obtained for the corresponding linear combination of regression parameters. We report the estimates of Equation (4.4.1) in Table 4.A.2 of the Appendix. Following our experimental design, we start with the analysis of choices under risk and then turn the intertemporal decisions.

4.4.2 Choice under risk

We can characterize the shape of the flow utility function up to the third derivative from the subjects' choices under risk.

Figure 4.4.1 shows the cumulative distribution of the required compensation payments in the risk apportionment tasks. This non-parametric analysis yields two main findings, which we discuss in turn.

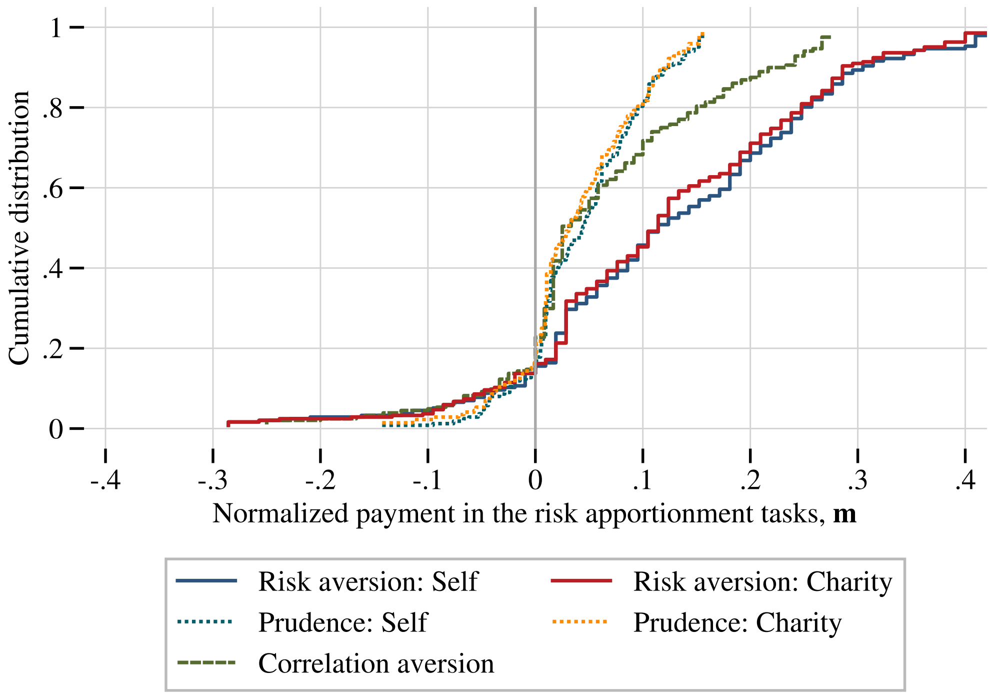

Figure 4.4.1. Risk preferences

Notes: This figure plots the cumulative distribution function of the normalized compensation payments *m* for each of the five stages of the risk apportionment tasks. For each risky choice, we first divide the indifference points by the expected value of the corresponding base lottery without compensation to render choices comparable (see Table 4.3.1 for an overview of each stage). For each stage, we then obtain *m* by taking the average of the three normalized lottery choices. The figure then plots the cumulative distribution function of *m* for each stage (N = 244). "Risk aversion: Self" and "Risk aversion: Charity" show the distribution of second-order risk attitudes over self-euros and charity-euros. "Prudence: Self" and "Prudence: Charity" show the distribution of third-order risk attributes over self-euros and charity-euros. "Correlation aversion" shows the distribution of the multivariate risk aversion over self-euros and charity-euros.

More than 80% of subjects display second- and third-order risk aversion for self-euros and charity-euros. We can neither reject the null hypothesis that people are on average *equally* risk-averse in both domains (paired Wilcoxon signed-rank test, $p = 0.251$) nor that risk preferences in both domains are equally distributed (Kolmogorov–Smirnov test, $p = 0.786$). In the following, we will thus assume that the single-attribute utility functions representing utility from self-euros and charity-euros only differ by a multiplicative constant.[18] We also observe a strong positive correlation ($\rho = 0.671$) between subjects' third-order risk aversion (prudence) in the self- and other domain.

18. The most commonly used one- and two-parameter families of utility functions are pinned down (up to a linear transformation) by their second- and third-order risk aversion.

Result 4. *Subjects exhibit highly similar attitudes towards risk in payments of self-euros and charity-euros. This observation implies that the corresponding single-attribute utility functions have equal curvature.*

We classify more than 80% of subjects as correlation averse. The risk apportionment tasks deliver a non-parametric measure of the condition for correlation aversion, namely that the cross-derivative with respect to payments in self-euros and charity-euros is negative.

Result 5. *Subjects overwhelmingly display correlation aversion. This implies that the multi-attribute utility function is not additively separable,* $u(w,g) \neq f(w) + h(g)$.

Summing up, we document the non-separability of multi-attribute utility and identical curvatures of the single-attribute utility functions. Both features inform our analysis of intertemporal choices from now on.

4.4.3 Intertemporal choice

We test our earlier predictions to disentangle consequence-dated and choice-dated prosocial utility. Our earlier finding that the single-attribute utility functions have the same curvature allows us to derive slightly more general conclusions than under the nested case of linear utility.

We start with the univariate discounting tasks (stages *UD-CHARITY* and *UD-SELF*). Here, subjects only face a time trade-off, but no trade-off across domains. Figure 4.4.2 shows the net present values of delayed payments of self-euros and charity-euros. We plot the average stated amounts for the subjective evaluation in self-euros m^S for a payment of one self-euro that is delayed by an amount of time τ. We report the same result for the subjective evaluation in charity-euros m^C for a donation of one euro that is also delayed by τ.

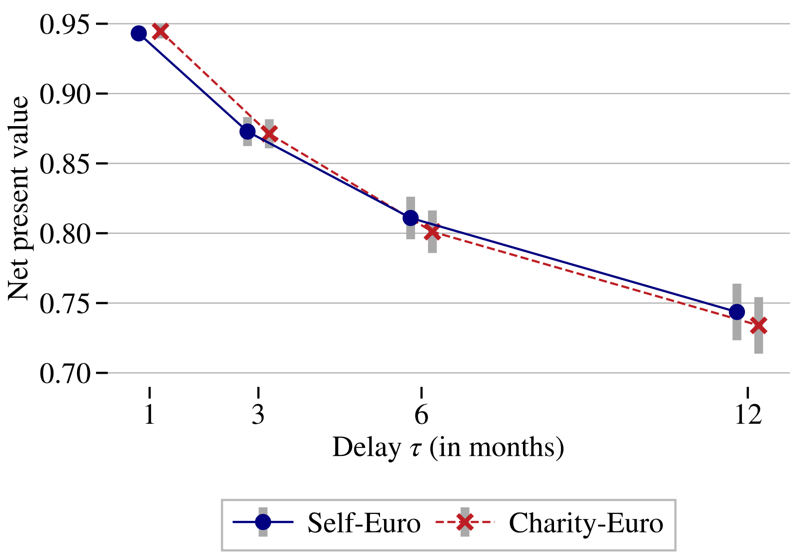

Figure 4.4.2. Univariate discounting

Notes: This figure displays the net present value of a dated payment of one self-euro (blue markers) and the net present value of one charity-euro (red markers) with different delays (N = 244). The net present values are calculated from choices between smaller-sooner and larger-later payments to the subjects or donations. 95% confidence intervals of the mean are calculated according to Morey (2008) and Cousineau (2005).

The net present values are identical between the two domains and decreasing with time. The average stated amounts for m^S and m^C are statistically indistinguishable for all delays τ (paired Wilcoxon signed-rank test, $p > 0.58$ for any τ). This result has two implications related to Prediction 3. First, a decreasing net present value for delayed charity-euro payments is incompatible with the pure choice-dated model. If prosocial utility flows are entirely choice-dated, then delays in implementing the donation payment are simply irrelevant. Our finding is in line with consequence-dated prosocial utility flows. Second, and more compellingly, the discounted utility version of consequence-dated utility can accommodate identical net present values for delayed self-euros and charity-euros and identical curvatures of the single-attribute utility functions (Result 4). This suggests that the same discount factors, $D(\tau)$, are applied to future utility from self-euros and charity-euros. We can rule out the alternative explanation that there are separate discount factors for each domain as, established in Result 4, the univariate utility functions for self-euros and charity-euros have the same curvatures.

Result 6. *In univariate discounting tasks, net present values for delayed self-euro and charity-euro payments are identical and decreasing in the delay. Both patterns are consistent with consequence-dated, but not with choice-dated, prosocial utility.*

We now turn to the choice tasks designed to determine subjective exchange rates between self-euros and charity-euros for different delays (stage *ER*). In these tasks, subjects only face a cross-attribute trade-off, since the time of the payments are the same. Figure 4.4.3 shows the average subjective exchange rates F_τ between contemporaneous self-euros and charity-euros.

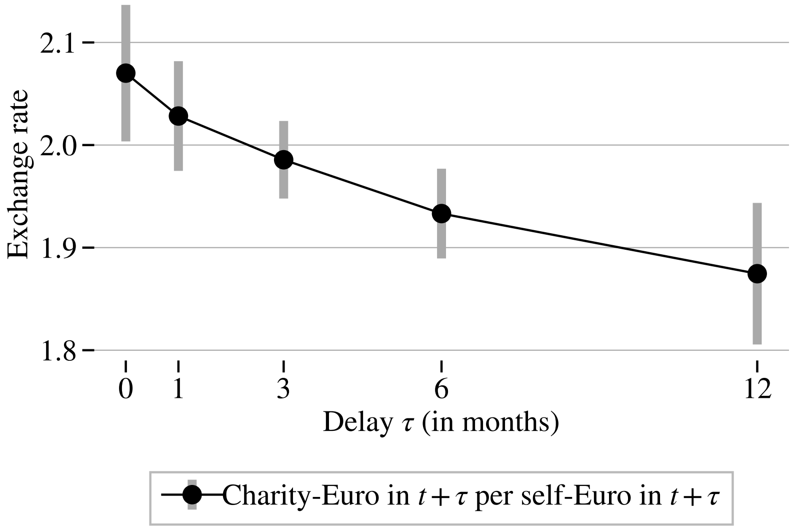

Figure 4.4.3. Forward exchange rate

Notes: This figure displays the estimated subjective exchange rates between contemporaneous payments to the subjects and donations, i.e. the number of charity-euros per contemporaneous self-euro. Note that "0 months" indicates payments initiated after three days. 95% confidence intervals of the mean are calculated according to Morey (2008) and Cousineau (2005).

The level of the subjective exchange rate is always above one indicating that subjects on average prefer payments to themselves over equally sized donations (paired t tests at each delay, $p < 0.001$). For the earliest payment date of only three days, subjects exhibit an exchange rate of approximately $F_\tau = 2.07$. One self-euro is valued about twice as much as one charity-euro. More strikingly, we find that the valuation of a self-euro per contemporaneous charity-euro decreases in the delay τ (paired t tests for the change in delay τ relative to base period, $p_1 = 0.245$, $p_3 = 0.031$, $p_6 = 0.003$, $p_{12} < 0.001$). This means that when the common delay of two payments – one denominated in self-euros and one in charity-euros – increases, our subjects develop a relative preference for charity-euros. Put differently, in these types of choices that only involve the same delay τ in both domains, subjects discount self-euros faster than charity-euros.

A declining forward exchange rate has two implications regarding Prediction 4. First, we cannot rationalize this pattern with the discounted utility version of consequence-dated prosocial utility. If we apply the same discount function to self-euros and charity-euros, the discount factors cancel out as the delays in the two payments are the same. Second, this finding is compatible with choice-dated prosocial utility. If delayed self-euro payments generate delayed utility flows that are discounted, but delayed donations are only associated with choice-dated utility flows, an increase in the common delay affects the discounted utility from self-euros, while leaving the utility derived from donations unaffected. Note that we do not have to invoke the shape of the utility function for this argument: the exchange rate finding is incompatible with discounted utility irrespective of utility curvatures.

Result 7. *Subjective exchange rates between self-euros and charity-euros are declining over time, i.e. a common delay makes self-euros relatively less valuable than charity-euros. This pattern is explained by choice-dated, but not by consequence-dated, prosocial utility.*

Finally, recall that choice tasks on multi-attribute discounting (stages *MD – CHARITY* and *MD – SELF*) combine a cross-attribute trade-off with a time trade-off within a single decision. Our participants had to decide what amount in one domain payable at a later date would make them indifferent to a given amount in the other domain payable at an earlier date. Figure 4.4.4 shows the average net present value of one delayed self-euro when expressed in charity-euros today and vice versa.

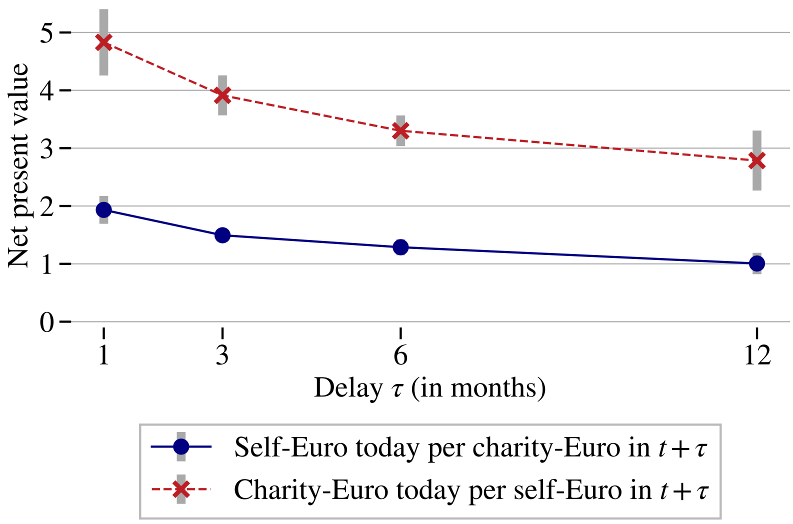

Figure 4.4.4. Multivariate discounting

Notes: This figure displays estimated net present values of delayed payments (N = 244). Red markers indicate the net present value of one delayed self-euro expressed in charity-euros today. Blue markers indicate the net present value of one delayed charity-euro expressed in self-euros today. 95% confidence intervals of the mean are calculated according to Morey (2008) and Cousineau (2005).

There are three distinct patterns in our data. First, subjects on average demand less compensation in self-euros at the earlier date for giving up a donation at a later date than vice versa (paired t test for each delay, $p < 0.01$).[19] Intuitively, given that

19. Specifically, the average WTA for giving up self-euros today for charity-euros tomorrow (WTA^{sc}_{τ}) is higher than the average WTA for giving up charity-euros today for self-euros tomorrow (WTA^{cs}_{τ}). In particular, we have $\min_{\tau} \text{WTA}^{sc}_{\tau} > \max_{\tau} \text{WTA}^{cs}_{\tau}$.

subjects value one self-euro roughly twice as much as a contemporaneous charity-euro, they will require less compensation in their preferred category (self-euros) than in the inferior category (charity-euros). Second, the net present values decrease in the delay of the later payment, implying that payments of both self-euros and charity-euros are valued less as their delay increases (paired t-tests between adjacent delays, $p < 0.01$). Third, we find that the net present value of delayed charity-euros decreases less quickly in the delay τ than the net present value of delayed self-euros (paired t-tests for the difference in rates of change for compensations in self-euros and charity-euros for each time difference, $p < 0.01$).

These non-parametric results relate to Prediction (5) as follows: a decreasing net present value of delayed charity-euros is at odds with pure choice-dated prosocial utility, as the payment date of charity-euros should be inconsequential in that case. However, all three patterns are compatible with consequence-dated prosocial utility. A decreasing net present value of more delayed donations naturally follows from stronger discounting. The level differences as well as the difference in slopes are predicted by a lower marginal utility from charity-euros.

Result 8. *In cross-attribute intertemporal decisions, net present values of delayed charity-euro payments are decreasing in their delay, and they are lower and decrease less quickly than required rates of return on delayed self-euros. These patterns are explained by consequence-dated, but not by choice-dated, prosocial utility.*

In summary, our reduced-form analyses provide strong evidence for the existence of *both* choice-dated and consequence-dated components of prosocial utility. We now develop and estimate a structural model that reproduces the documented patterns with a single set of preferences.

4.5 Structural estimation

Our following structural analysis has two objectives (DellaVigna, 2018). First, we assess the ability of our proposed model of intertemporal prosocial utility to generate the qualitative reduced-form patterns with a quantitatively reasonable parameterization. Second, the estimated model allows us to assess the relative importance of choice-dated and consequence-dated prosocial utility.

We first outline and motivate the functional form of our utility function, provide details about our estimation routine, and discuss the results from a representative agent model before we turn to an individual-level estimation.

4.5.1 Setup

Building on our conceptual framework and reduced-form results, we posit the following parametric form for our intertemporal utility function:

$$U_t = \underbrace{\alpha \mathbb{1}\left(\sum_{\tau=0}^{\infty} g_{t+\tau} > 0\right)}_{\text{chioce-dated}} + \underbrace{\sum_{\tau=0}^{\infty} \delta^{\tau}\left(w x_{t+\tau}^{\beta} + (1-w) g_{t+\tau}^{\beta}\right)}_{\text{consequence-dated}}. \tag{4.5.1}$$

The first part represents choice-dated prosocial utility, while the second part captures consequence-dated utility. In this parameterization, α is the choice-dated prosocial utility derived from donating, and δ denotes the one-month utility discount factor. We capture pure altruism by $1-w$, as it describes the relative value of one charity-euro to a current self-euro. $1-\beta$ refers to the coefficient of univariate relative risk aversion.

The key elements of our specification follow our reduced-form analysis and the existing literature.[20] First, our earlier findings suggest that we include both choice- and consequence-dated utility. Second, for the flow utility function, we document in our reduced-form analysis that the curvature of the univariate utility from self-euros and the univariate utility from charity-euros have the same curvature and only differ in scale. We, therefore, assume a common parameter, β, to capture the curvature of the utility function when choices involve only one recipient. Third, we also find strong evidence of multivariate risk aversion in our reduced-form analysis, implying a non-additively separable flow utility function in the consequence-dated utility component. Finally, we assume standard exponential discounting as our data only includes payment dates in the future, allowing us to abstract from present bias and to economize on parameters in our baseline specification.

We drop a small number of questions and individuals from our estimation sample. While it is possible to explicitly incorporate a parameter of correlation aversion in the functional form, our primary focus is on intertemporal prosocial utility. Indeed, correlation aversion should only affect 3 out of 36 choices. As such, we abstain from modeling correlation aversion and exclude all choices from stage *X – RA* in our estimation. In addition, some subjects display a very high degree of risk aversion in the stages *RA – SELF* and *RA – CHARITY*. As highlighted in Wakker (2008), a CRRA utility function has difficulties matching this behavior, as a constant relative risk aversion greater than one is outside the theoretical range of our structural model. Thus, we exclude 44 subjects with an average normalized switching point greater than 0.9 in the stages *RA – SELF* and *RA – CHARITY* to avoid corner solutions.

4.5.2 Estimation

The experiment is carefully designed to provide the required variation to jointly identify the four parameters $\theta = (\alpha, \beta, \delta, w)$ in Equation (4.5.1). Univariate risk aversion, $1-\beta$, is identified from the risky choices in Part B. Conditional on $1-\beta$,

20. Andreoni and Miller (2002), Andersen et al. (2018), and Fisman, Kariv, and Markovits (2007) use a similar functional form.

the discount factor δ is separately identified from the univariate discounting stage in Part A of the experiment. The subjective exchange rate from stage ER provides identifying variation for the choice-dated prosocial utility parameter, α. We identify the pure altruism parameter, $1-w$, from choices involving trade-offs between self-euros and charity-euros such as stage *MD – SELF*, *MD – CHARITY*, and *ER*.

We estimate the structural parameters of our model using a minimum-distance estimator (Newey and McFadden, 1994). Let $m(\theta)$ denote the moments predicted by our structural model, and $\hat{m}$ the vector of observed moments. The minimum-distance estimator selects the parameters $\hat{\theta}$ that minimize the distance the squared distance between the observed and predicted moments. The estimates $\hat{\theta}$ are defined by:

$$\hat{\theta} \;=\; \underset{\theta}{\operatorname{argmin}} \quad (m(\theta)-\hat{m})'W(m(\theta)-\hat{m}), \tag{4.5.2}$$

where W is a positive definite weighting matrix. We chose a minimum-distance estimator over a maximum likelihood estimator, because it is more robust to outliers that are unlikely according to the model. This challenge is particularly prevalent in the context of charitable giving (DellaVigna, List, and Malmendier, 2012; DellaVigna, 2018).

As a vector of moments $\hat{m}$, we use the average normalized switching point in each of our remaining 33 price lists. We normalize individual switching points by applying a linear transformation that maps each price list onto the unit interval such that $\hat{m} \in [0,1]^{33}$. For the choice of the weighting matrix, we follow the suggestion by Altonji and Segal (1996) and use the diagonal of the inverse of the variance-covariance matrix of our empirical moments. We provide additional details about the implementation and reliability of our estimation approach in Appendix 4.B.

4.5.3 Results

We estimate two models to learn about our parameters of interest. We first estimate a representative agent model that rules out any parameter heterogeneity. Then, we leverage the rich within-subject variation of our data. We estimate the utility function at the subject level and obtain estimated preferences $\hat{\theta}_i$ for each subject (Fisman, Kariv, and Markovits, 2007; Augenblick and Rabin, 2018).

First, we consider the representative agent model. Figure 4.5.1 displays the point estimates and the corresponding 95% confidence intervals for the model parameters.

Our estimated parameter values are all reasonable and, where applicable, in line with the existing literature. For example, we estimate a one-month discount factor of 0.991, which corresponds to a one-year discount factor of 0.906, similar to results observed by Andersen et al. (2018). We estimate a univariate relative risk aversion parameter of 0.808, and we find evidence for a consequence-dated prosocial utility

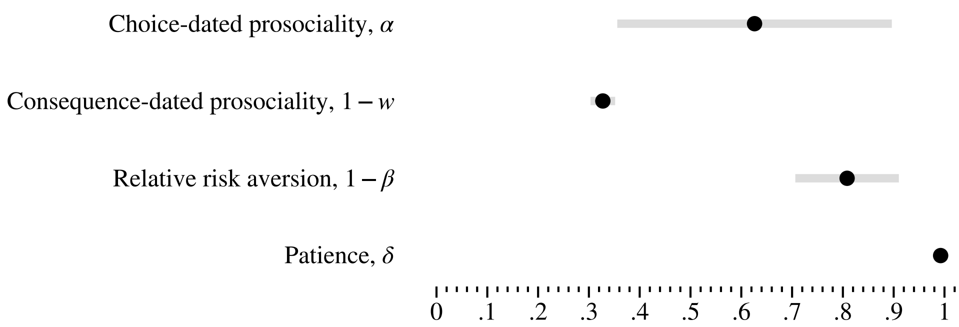

Figure 4.5.1. Parameter estimates: Representative agent

Notes: This figures displays the point estimates (black marker) and 95% confidence intervals (gray lines) of the representative agent parameter estimation (N = 200). α is the marginal choice-dated prosocial utility from giving. $1-w$ is the weight on utility from charity-euros in the stationary flow utility function. $1-\beta$ is the coefficient of univariate relative risk aversion. δ is the one-month discount factor.

component. Our point estimate of $1-\hat{w}=0.32$ implies that a donation of 50 euros provides roughly half (i.e. $\frac{1-w}{w}$) of the utility of an identically-dated 50-euro payment to the subject. This magnitude is consistent with our reduced-form estimate of the subjective exchange rate. In addition, there is a choice-dated utility component. We estimate a value of $\hat{\alpha}=0.62$ that implies that a donation of 50 euros in one month provides about 40% of the utility associated with a 50-euro payment to the subject with the same delay.

Next, we turn to the individual-level estimation to investigate the role of preference heterogeneity in our sample. We find considerable heterogeneity in preferences. Figure 4.5.2 shows the marginal distribution of each preference parameter. The median subject exhibits a consequence-dated prosociality parameter $1-\hat{w}$ of 0.353, which is in line with the estimate for a representative agent. At the same time, about 20% of respondents have parameter estimates $(1-\hat{w}=0)$ that suggest almost no concern for the consequences of their decisions for others. Slightly fewer than 60% of our subjects have parameter estimates $\hat{\alpha}>0$ that suggest the presence of choice-dated prosociality. Among this group, the degree of choice-dated motivation is widely dispersed with a median parameter estimate of 0.481.

Looking at the joint distribution between $\hat{\alpha}$ and $(1-\hat{w})$ we find a negative correlation of $\rho=-0.417$. This strong correlation suggests that the prosocial motivations underlying these differently-dated utility flows are substitutes rather than complements at the individual level. Put differently, choice-dated and consequence-dated motivations characterize different types of people. Our data are compatible with the interpretation that, while some people donate out of pure altruism, others are driven by the feeling of warm glow.

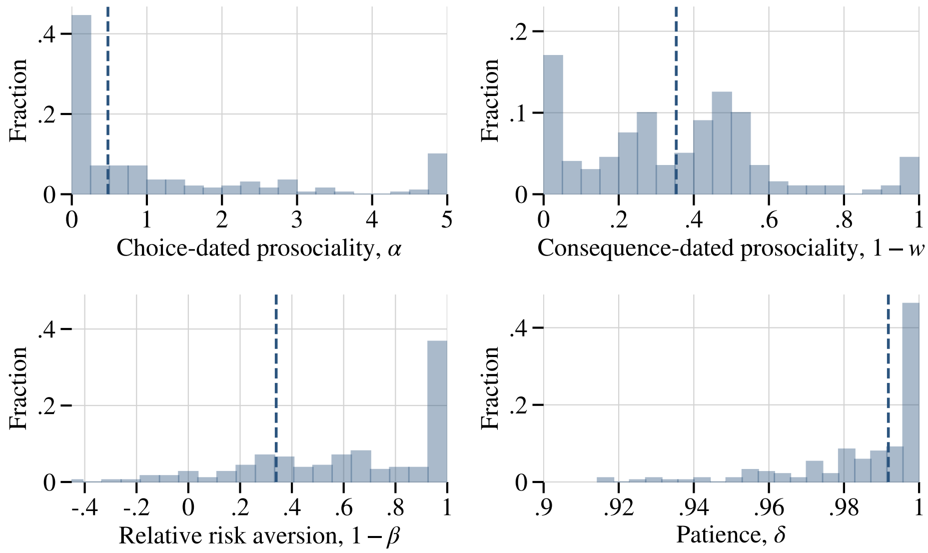

Figure 4.5.2. Parameter estimates: Individual-level estimates

Notes: This figure plots the marginal distribution of the model parameters at the subject-level (N = 200). It shows the fraction of the sample that is contained in each bin. The dashed vertical line indicates the median of the distribution. The distribution of $1 - \beta$ excludes fifteen subjects with a coefficient of relative risk aversion smaller than -0.50. The distribution of δ excludes twelve subjects with a one-month discount factor below 0.90.

The declining subjective exchange rate between self-euro and charity-euro payments with the same delay is one of the core findings from our reduced form analysis (see Appendix Figure 4.A.3). Our estimated model replicates this pattern, and we discuss the implications for the median of the individual-specific parameters. First, consider two payments executed in a month from today. A 50-euro payment to the subject in one month provides 2.42 utils of discounted utility to a subject, whereas a 50-euro donation in one month provides 1.32 utils from consequence-dated utility flows and 0.48 utils from the choice-dated utility flow. Second, consider two payments executed in a year from today. A 50-euro payment to the subject in a year provides 2.21 discounted utils to a subject, whereas a 50-euro donation in a year provides only 1.21 utils from consequence-dated utility flows and still 0.48 utils from the choice-dated utility flow. Going from a delay of only a month to a full year leads to an increase in the relative weight of the choice-dated utility prosocial utility from 26.7% to 28.5% in this example. As a consequence, the forward exchange rate decreases by 0.28 euros when payments are executed in a year from today rather than a month. This change is remarkably close to our observed decrease of 0.20 euros in our experiment.

4.6 Conclusion

We study the intertemporal dimension of prosocial behavior and propose a distinction between choice-dated and consequence-dated flows of prosocial utility. This conceptual approach generalizes differences between psychological motivations explored in the existing literature and delivers testable implications for intertemporal prosocial behavior. Empirically, we conduct a high-stakes donation experiment that provides a comprehensive characterization of the intertemporal multi-attribute utility function using reduced-form and structural approaches. We find that the majority of individuals exhibit both choice-dated and consequence-dated prosocial utility. Furthermore, both motives are quantitatively meaningful, and there is a strong negative correlation between their importance as individuals are either primarily motivated by choice-dated or consequence-dated considerations.

We conclude with three comments on the limitations and potential promise of the approach taken in this paper. First, the proposed conceptual distinction between consequence-dated and choice-dated utility is deliberately chosen to bridge theoretical work on intertemporal choice with largely empirical work on specific prosocial motivations such as warm glow and pure altruism. At the same time, this taxonomy remains a reduced-form perspective on the psychological mechanisms underlying prosocial behavior. It is thus complementary to work that sheds light on the sources of pure time preferences about the outcomes of others. For example, our approach and findings provide a motivation to further examine *why* people prefer helping others sooner rather than later. Second, we abstract from the implications of our approach for time-inconsistent behavior. This topic has received significant attention following work on present-biased preferences and is the focus of related work. Third, while the present paper introduces a toolkit for analyzing the time structure of prosocial utility flows and hints at the usefulness of this approach for understanding prosocial decision-making, it does so in a specific high-stakes donation context using a specific experimental paradigm that relies on the well-studied multiple price list methodology and monetary payments. One avenue for future work is to examine the implications of whether and how intertemporal prosocial motivations interact with these design choices.

References

Altonji, Joseph G., and Lewis M. Segal. 1996. "Small-sample bias in GMM estimation of covariance structures." *Journal of Business & Economic Statistics* 14 (3): 353–66. [290]

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström. 2008. "Eliciting risk and time preferences." *Econometrica* 76 (3): 583–618. [266, 277]

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström. 2018. "Multiattribute utility theory, intertemporal utility and correlation aversion." *International Economic Review* 59 (2): 537–55. [269, 289, 290]

Andreoni, James. 1989. "Giving with impure altruism: Applications to charity and ricardian equivalence." *Journal of Political Economy* 97 (6): 1447–58. [264, 268, 272, 305]

Andreoni, James. 1990. "Impure altruism and donations to public goods: A theory of warm-glow giving." *Economic Journal* 100 (401): 464–77. [264, 272]

Andreoni, James. 1993. "An experimental test of the public-goods crowding-out hypothesis." *American Economic Review* 83 (5): 1317–27. [268]

Andreoni, James, Michael A. Kuhn, and Charles Sprenger. 2015. "Measuring time preferences: A comparison of experimental methods." *Journal of Economic Behavior & Organization* 116: 451–64. [277]

Andreoni, James, and John Miller. 2002. "Giving according to GARP: an experimental test of the consistency of preferences for altruism." *Econometrica* 70 (2): 737–53. [289]

Andreoni, James, and Marta Serra-Garcia. 2021. "Time inconsistent charitable giving." *Journal of Public Economics* 198: 104391. [265, 268]

Andreoni, James, and Charles Sprenger. 2012. "Estimating time preferences from convex budgets." *American Economic Review* 102 (7): 3333–56. [266, 277]

Ariely, Dan, and Michael I. Norton. 2009. "Conceptual consumption." *Annual Review of Psychology* 60: 475–99. [269]

Attema, Arthur E., Han Bleichrodt, Yu Gao, Zhenxing Huang, and Peter P. Wakker. 2016. "Measuring discounting without measuring utility." *American Economic Review* 106 (6): 1476–94. [277]

Augenblick, Ned. 2019. "Short-term time discounting of unpleasant tasks." *Working Paper, UC Berkeley*, [266, 276]

Augenblick, Ned, and Matthew Rabin. 2018. "An experiment on time preference and misprediction in unpleasant tasks." *Review of Economic Studies* 86 (3): 941–75. [266, 290]

Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela. 2020. "How soon is now? Evidence of present bias from convex time budget experiments." *Experimental Economics* 23 (2): 294–321. [266, 270, 276]

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and prosocial behavior." *American Economic Review* 96 (5): 1652–78. [264, 269]

Bolton, Gary, and Elena Katok. 1998. "An experimental test of the crowding out hypothesis: The nature of beneficent behavior." *Journal of Economic Behavior & Organization* 37 (3): 315–31. [268]

Bolton, Gary E., and Axel Ockenfels. 2000. "ERC: A theory of equity, reciprocity, and competition." *American Economic Review* 90 (1): 166–93. DOI: 10.2307/117286. [264]

Breman, Anna. 2011. "Give more tomorrow: Two field experiments on altruism and intertemporal choice." *Journal of Public Economics* 95 (11): 1349–57. [265, 268]

Bruner, David M. 2011. "Multiple switching behaviour in multiple price lists." *Applied Economics Letters* 18 (5): 417–20. [277]

Byrd, Richard H., Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. 1995. “A limited memory algorithm for bound constrained optimization.” *SIAM Journal on Scientific Computing* 16 (5): 1190–208. [304]

Chabris, Christopher F., David Laibson, Carrie Morris, Jonathon Schuldt, and Dmitry Taubinsky. 2008. “Individual laboratory-measured discount rates predict field behavior.” *Journal of Risk and Uncertainty* 37 (2-3): 237–69. [266]

Chakraborty, Anujit, Evan M. Calford, Guidon Fenig, and Yoram Halevy. 2017. “External and internal consistency of choices made in convex time budgets.” *Experimental Economics* 20 (3): 687–706. [266, 277]

Chapman, Gretchen B. 1996. “Temporal discounting and utility for health and money.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22 (3): 771. [269]

Charness, Gary, and Matthew Rabin. 2002. “Understanding social preferences with simple tests.” *Quarterly Journal of Economics* 117 (3): 817–69. [264]

Chen, Daniel L., Martin Schonger, and Chris Wickens. 2016. “oTree—An open-source platform for laboratory, online, and field experiments.” *Journal of Behavioral and Experimental Finance* 9: 88–97. [281]

Chopra, Felix, Philipp Eisenhauer, Armin Falk, and Thomas Graeber. 2021. “Intertemporal Altruism.” *IZA Discussion Paper No. 14059*, 89. DOI: http://dx.doi.org/10.2139/ssrn.3771742.

Cohen, Jonathan D., Keith Ericson, David Laibson, and John Miles White. 2019. “Measuring time preferences.” *Journal of Economic Literature*, [266, 271, 276]

Cousineau, Denis. 2005. “Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson’s method.” *Tutorials in Quantitative Methods for Psychology* 1 (1): 42–45. [282, 285–287]

Cubitt, Robin, Rebecca McDonald, and Daniel Read. 2018. “Time matters less when outcomes differ: Unimodal vs. Cross-modal comparisons in intertemporal choice.” *Management Science* 64 (2): 873–87. [269]

Cubitt, Robin, and Daniel Read. 2007. “Can intertemporal choice experiments elicit preferences for consumption.” *Experimental Economics* 10 (4): 369–89. [266]

DellaVigna, Stefano. 2018. *Structural Behavioral Economics*. Vol. 1, Elsevier B.V., 613–723. DOI: 10.1016/bs.hesbe.2018.07.005. [265, 288, 290]

DellaVigna, Stefano, John List, and Ulrike Malmendier. 2012. “Testing for altruism and social pressure in charitable giving.” *Quarterly Journal of Economics* 127 (1): 1–56. [290]

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. 2017. “The robustness and pervasiveness of sub-additivity in intertemporal choice.” *Working Paper*, [277]

Dreber, Anna, Drew Fudenberg, David K. Levine, and David G. Rand. 2016. “Self-control, social preferences and the effect of delayed payments.” *SSRN Working Paper*, [265, 268]

Dufwenberg, Martin, and Georg Kirchsteiger. 2004. “A theory of sequential reciprocity.” *Games and Economic Behavior* 47 (2): 268–98. DOI: 10.1016/j.geb.2003.06.003. [264]

Ebert, Sebastian, and Gijs van de Kuilen. 2015. “Measuring multivariate risk preferences.” *SSRN Working Paper*, [279]

Ebert, Sebastian, and Daniel Wiesen. 2014. “Joint measurement of risk aversion, prudence, and temperance.” *Journal of Risk and Uncertainty* 48 (3): 231–52. [280]

Eeckhoudt, Louis, and Harris Schlesinger. 2006. “Putting risk in its proper place.” *American Economic Review* 96 (1): 280–89. [279]

Epstein, Larry G., and Stephen M. Tanny. 1980. “Increasing generalized correlation: A definition and some economic consequences.” *Canadian Journal of Economics* 13 (1): 16–34. [279]

Ericson, Keith Marzilli, and Jawwad Noor. 2015. "Delay functions as the foundation of time preference: Testing for separable discounted utility." *NBER Working Paper*, [277]

Falk, Armin, and Urs Fischbacher. 2006. "A theory of reciprocity." *Games and Economic Behavior* 54 (2): 293–315. DOI: 10.1016/j.geb.2005.03.001. [264, 269]

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A theory of fairness, competition, and cooperation." *Quarterly Journal of Economics* 114 (3): 817–68. [264]

Fisman, Raymond, Shachar Kariv, and Daniel Markovits. 2007. "Individual preferences for giving." *American Economic Review* 97 (5): 1858–76. [289, 290]

Frederick, Shane. 2006. "Valuing future life and future lives: A framework for understanding discounting." *Journal of Economic Psychology* 27 (5): 667–80. [269]

Gabler, Janos. 2020. "A python tool for the estimation of (structural) econometric models." https://github.com/OpenSourceEconomics/estimagic. [304]

Halevy, Yoram. 2015. "Time consistency: Stationarity and time invariance." *Econometrica* 83 (1): 335–52. [266, 270, 276]

Hardisty, David J., and Elke U. Weber. 2009. "Discounting future green: Money versus the environment." *Journal of Experimental Psychology: General* 138 (3): 329. [269]

Holt, Charles A., and Susan K. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review* 92 (5): 1644–55. [277]

Kim, Hyunji, Simone Schnall, and Mathew P. White. 2013. "Similar psychological distance reduces temporal discounting." *Personality and Social Psychology Bulletin* 39 (8): 1005–16. [269]

Kolappan, C., R. Subramani, V. Kumaraswami, T. Santha, and P. R. Narayanan. 2008. "Excess mortality and risk factors for mortality among a cohort of TB patients from rural south India." *International Journal of Tuberculosis and Lung Disease* 12 (1): 81–86. [275]

Kölle, Felix, and Lukas Wenner. 2018. "Time-inconsistent generosity: Present bias across individual and social contexts." *Working Paper*, [265, 268]

Konow, James. 2010. "Mixed feelings: Theories of and evidence on giving." *Journal of Public Economics* 94 (3-4): 279–97. [268]

Kovarik, Jaromir. 2009. "Giving it now or later: Altruism and discounting." *Economics Letters* 102 (3): 152–54. [265, 268]

McClure, Samuel M., Keith M. Ericson, David I. Laibson, George Loewenstein, and Jonathan D. Cohen. 2007. "Time discounting for primary rewards." *Journal of Neuroscience* 27 (21): 5796–804. [269]

Meer, Jonathan, and Benjamin A. Priday. 2020. "Generosity across the income and wealth distributions." *NBER Working Paper*, [268]

Montiel Olea, José Luis, and Tomasz Strzalecki. 2014. "Axiomatization and measurement of quasi-hyperbolic discounting." *Quarterly Journal of Economics* 129 (3): 1449–99. [277]

Morey, Richard. 2008. "Confidence intervals from normalized data: A correction to Cousineau (2005)." *Tutorials in Quantitative Methods for Psychology* 4 (2): 61–64. [282, 285–287]

Newey, Whitney K., and Daniel McFadden. 1994. "Large sample estimation and hypothesis testing." In *Handbook of Econometrics*. Edited by R. F. Engle and D. L. McFadden. Vol. 4, Elsevier, 2111–245. [290]

Richard, Scott F. 1975. "Multivariate risk aversion, utility independence and separable utility functions." *Management Science* 22 (1): 12–21. [279]

Rogers, Todd, and Max H. Bazerman. 2008. "Future lock-in: Future implementation increases selection of 'should' choices." *Organizational Behavior and Human Decision Processes* 106 (1): 1–20. [265, 268]

Schelling, T. C. 1988. "The mind as a consuming organ." In *Decision Making: Descriptive, Normative, and Prescriptive Interactions*. Edited by David E. Bell, Howard Raiffa, and Amos Tversky. Cambridge University Press, 343–57. [269]

Schubert, Renate, Martin Brown, Matthias Gysler, and Hans Wolfgang Brachinger. 1999. "Financial decision-making: Are women really more risk-averse?" *American Economic Review* 89 (2): 381–85. [277]

Starmer, Chris. 2000. "Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk." *Journal of Economic Literature* 38 (2): 332–82. [279]

Straetemans, Masja, Philippe Glaziou, Ana L. Bierrenbach, Charalambos Sismanidis, and Marieke J. van der Werf. 2011. "Assessing tuberculosis case fatality ratio: A meta-analysis." *PLOS ONE* 6 (6): 1–13. [275]

Tiemersma, Edine W., Marieke J. van der Werf, Martien W. Borgdorff, Brian G. Williams, and Nico J. D. Nagelkerke. 2011. "Natural history of tuberculosis: Duration and fatality of untreated pulmonary tuberculosis in HIV negative patients: a systematic review." *PLOS ONE* 6 (4): 1–13. [275]

Wakker, Peter P. 2008. "Explaining the characteristics of the power (CRRA) utility family." *Health Economics* 17 (12): 1329–44. [289]

World Health Organization. 2020. "Tuberculosis (TB) Fact Sheet." URL: https://www.who.int/en/news-room/fact-sheets/detail/tuberculosis. [275]

Appendix 4.A Tables and Figures

Table 4.A.1. Summary statistics

	Observations	Mean	Std. dev.	Min	25th	Median	75th	Max
Age	244	25	5.5	18	22	23	26	61
Female	244	.57	.5	0	0	1	1	1
Household income	244	1,446	1,133	0	650	1,000	2,000	4,000
Savings	244	.54	.5	0	0	1	1	1
Education (years)	244	16	3.5	3	15	16	18	29
Student	244	.91	.29	0	1	1	1	1
Political orientation	244	2.3	1.3	0	1	2	3	6
Siblings	244	1.5	1.2	0	1	1	2	7
Raven score	244	6.1	1.7	0	5	6	7	10

Notes: This table shows summary statistics for the full sample. "Household income" is the self-reported total monthly household income after taxes and transfers (in euros). "Savings" is a binary variable taking the value of 1 if the subject reported that she is able to save money each month. "Education (years)" are the subject's total years of education starting from primary school. "Student" is a binary variable taking value of 1 if the subject is enrolled at a university degree program. "Political orientation" is measured on a scale from 1 ("rather left") to 7 ("rather right"). "Siblings" are the total number of siblings. "Raven score" is the number of correctly solved Raven matrices out of ten.

Table 4.A.2. Regression analysis of intertemporal choices

	Univariate discounting			Multivariate discounting			Exchange rate
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
charity-euro	-0.005		0.001	2.277		2.897	
	(0.008)		(0.004)	(0.535)		(0.758)	
1 month							-0.042
							(0.036)
3 months		-0.072	-0.070		-0.678	-0.439	-0.084
		(0.004)	(0.005)		(0.158)	(0.109)	(0.039)
6 months		-0.138	-0.132		-1.087	-0.646	-0.137
		(0.006)	(0.008)		(0.193)	(0.158)	(0.045)
12 months		-0.205	-0.199		-1.485	-0.927	-0.195
		(0.009)	(0.011)		(0.250)	(0.181)	(0.054)
3 months × charity-euro			-0.003			-0.478	
			(0.006)			(0.320)	
6 months × charity-euro			-0.011			-0.883	
			(0.009)			(0.373)	
12 months × charity-euro			-0.011			-1.117	
			(0.013)			(0.471)	
Constant	0.843	0.944	0.943	1.430	3.381	1.933	2.070
	(0.004)	(0.004)	(0.005)	(0.268)	(0.140)	(0.308)	(0.030)
N	1952	1952	1952	1952	1952	1952	1220
R^2	0.386	0.620	0.621	0.428	0.404	0.437	0.921
Subject FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows pooled OLS regression estimates where the unit of observation are subject-choices. In columns 1–3, we include all choices from the two univariate discounting stages (UD-S, UD-C). The dependent variable is the net present value $y_{i,\tau,d}$ of the delayed payment, where i denotes the subject, τ the delay in months, and d is the numéraire of the payments (self-euros or charity-euros). Columns 4–6 include all choices from the two multivariate discounting stages (MD-S, MD-C). The dependent variable is the net present value $y_{i,\tau,d}$ of the delayed payment using the type d of the earlier payment (self-euros or charity-euros) as numéraire. In column 7, we include all choices from the exchange rate stage ER. The dependent variable is the implied (forward) exchange rate $y_{i,\tau}$ at different delays τ. "Charity-euro" is a binary indicator variable taking the value of 1 if the numéraire of the earlier payment are charity-euros. "τ month(s)" is a binary indicator variable taking the value of 1 if the later payment is received with a delay of τ month(s), where $\tau = 1$ month is the omitted category in columns 1–6 and "0 months" is the omitted category in column 7. All regressions include subject fixed effects for the 244 subjects. Standard errors are clustered at the subject level and shown in parentheses.

Treffen Sie jetzt Ihre Entscheidung

Bitte geben Sie für jede Zeile in der folgenden Tabelle an, ob Sie **Option A** oder **Option B** wählen.

Option A		Option B
50,00 € heute **selbst erhalten.**	○ ○	**0,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**10,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**21,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**31,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**42,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**52,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**63,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**73,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**84,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**94,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**105,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**115,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**126,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**136,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**147,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**157,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**168,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**178,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**189,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**199,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**210,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**220,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**231,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**241,50 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**252,00 €** in 6 Monaten **an Operation ASHA spenden.**
50,00 € heute **selbst erhalten.**	○ ○	**262,50 €** in 6 Monaten **an Operation ASHA spenden.**

Automatische Ausfüllhilfe: Damit Sie weniger klicken müssen, haben wir eine Ausfüllhilfe aktiviert, die automatisch Auswahlfelder für Sie ausfüllt.

Weiter

Figure 4.A.1. Decision screen: Intertemporal choices

Notes: This is an example of the decision screen as seen by subjects in stage *MD – SELF* of the intertemporal choice part of the experiment. The original instructions in German are shown. In each row, subjects indicate whether they prefer option A or option B by selecting the appropriate circle in each row. Option A on the left-hand side offers 50 self-euros today. Option B on the right-hand side offers increasing amounts of charity-euros from zero to 262.50 euros. The amount will be wired to *Operation ASHA* in six months. All price lists in the intertemporal choice part of our experiment are presented in this format. We vary only (i) the amount offered in option B, (ii) the timing of payments (both for option A and option B), and (iii) whether payments are denoted in self-euros or charity-euros. The decision screens are otherwise identical.

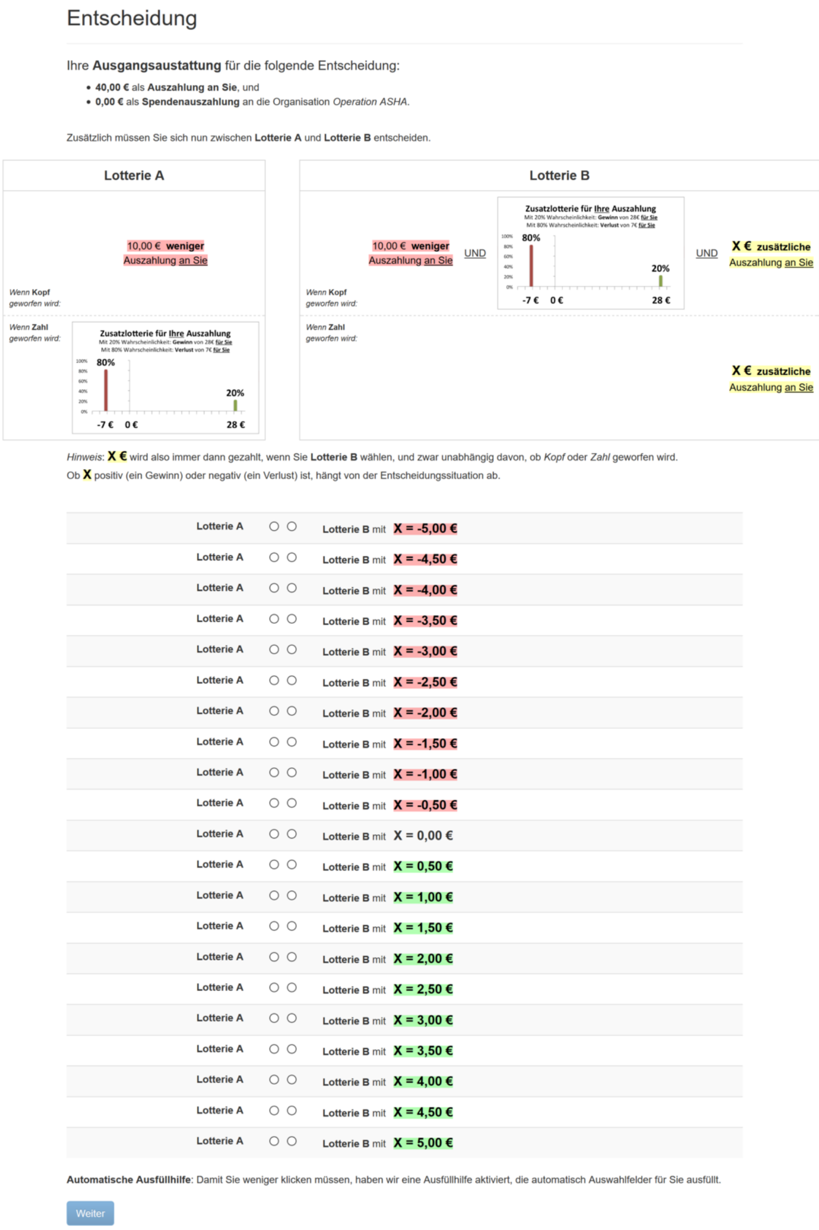

Figure 4.A.2. Decision screen: Risky lottery choices

Notes: This is an example of the decision screen as seen by subjects in stage *RA* – *SELF* of the risky choice part of the experiment. The original instructions in German are shown. At the top of the screen, subjects are informed about their initial endowment ***e*** of 40 self-euros and zero charity-euros. Next, subjects see two boxes that contain a visual representation of lottery A and lottery B. In each box, the upper part explains the consequences when the simulated coin toss yields head, whereas the lower part explains the consequences if it yields tails. In the lower part of the screen, subjects indicate whether they prefer lottery A or lottery B by selecting the appropriate circle in each row. The right-hand side shows the compensation amounts ***m*** that are to be added to lottery B. They range from -5.00 self-euros to 5.00 self-euros. All decisions in the risky choice part of our experiment are presented in this format. We vary only (i) the lotteries and (ii) the range of the compensation amounts. The decision screens are otherwise identical.

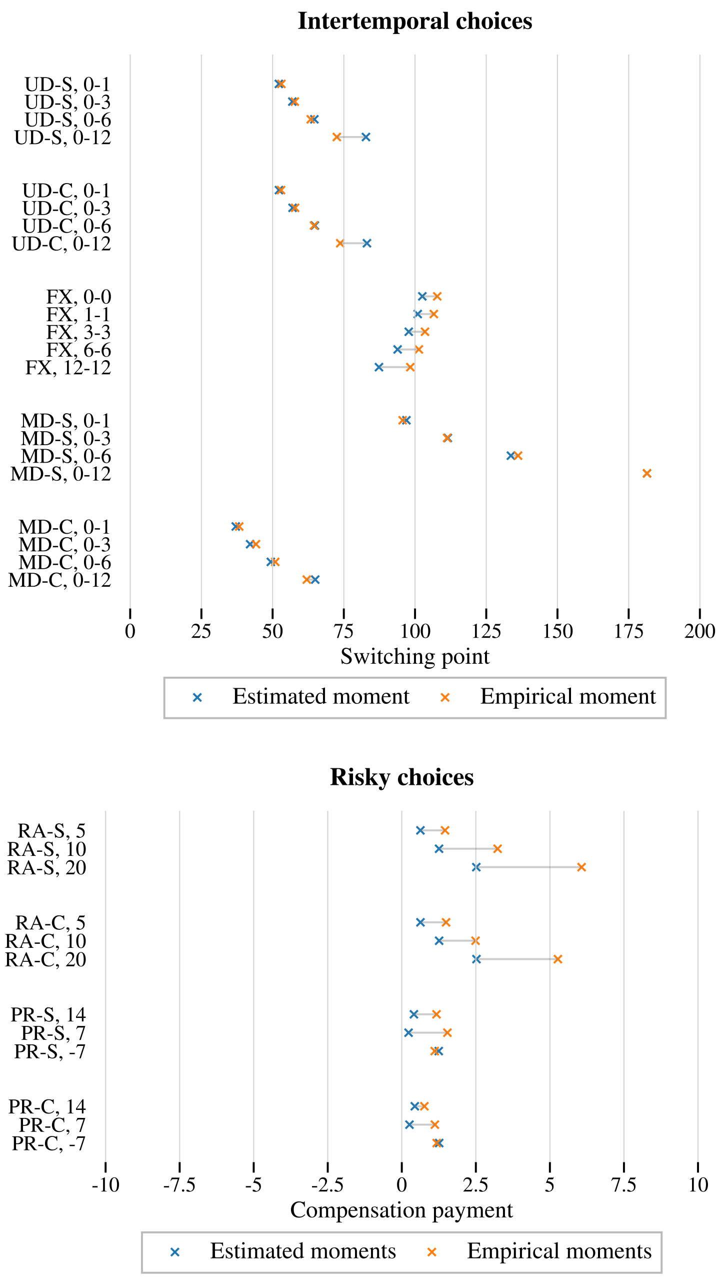

Figure 4.A.3. Structural estimation: Moments

Notes: This figure plots the empirical and the estimated moments for our estimation sample (N = 200). The moments are the average switching point in each of our 33 price lists. The upper panel shows moments for intertemporal choices, while the lower panel reports moments for risky choices from part B of the experiment. For intertemporal choices, labels on the vertical axis groups task by their stage (UD-S, UD-C, FX, MD-S, MD-C) and indicate the delay of the sooner and the later payment. For example, “6-6” means that both payments were made 6 months after the experiment. For risky choices, we indicate the size of the deduction R_2 (see Table 4.3.1 for more details).

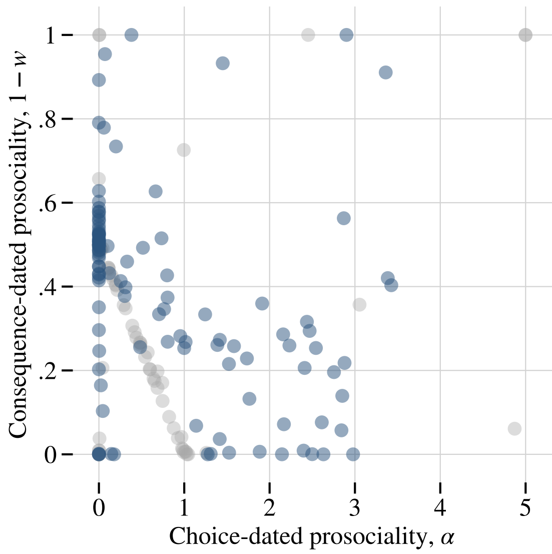

Figure 4.A.4. Structural model: Distribution of prosocial motivations

Notes: This figure shows the joint distribution (N = 200) of the choice-dated prosociality parameter, α, and the consequence-dated prosociality parameter, $1 - w$. The circles in dark gray indicate the subsample of subjects with a degree of risk aversion that is outside the range of the structural model, i.e. they have a coefficient of relative risk aversion greater than 0.90. The Spearman correlation is -0.417 in the full sample and -0.447 in the subsample.

Appendix 4.B Structural estimation

Practical estimation. To calculate the minimum-distance estimator $\hat{\theta}$, we employ the L-BFGS-B algorithm, which is appropriate for constrained optimization (Byrd, Lu, Nocedal, and Zhu, 1995).[21] We impose the following box constraints: $\delta \in (0,1]$ (positive discounting), $\beta \in [0,5]$, $\alpha \in [0,5]$ (non-negative choice-dated utility) and $w \in [0,1]$ (altruism weight between 0 and 1). As local minima are a natural concern in any structural estimation, we repeatedly estimate our model using ten randomly-chosen initial values from a uniform distribution over the parameter space. Moreover, we always include as initial values at least one parameter draw where $\alpha = 1 - w = 0$ to ensure that purely selfish preferences were in the consideration set of the estimator. As our final parameter estimate, $\hat{\theta}$, we choose the estimate with the minimum weighted distance among all ten estimates. We obtain standard errors from an estimator of the asymptotic variance-covariance matrix of the estimator:

$$(\hat{G}'W\hat{G})^{-1}(\hat{G}'W\hat{\Lambda}W\hat{G})(\hat{G}'W\hat{G})^{-1}, \tag{4.B.1}$$

where $\hat{G} = N^{-1}\sum_{i=1}^{N} \nabla_{\theta} m_i(\hat{\theta})$ and $\hat{\Lambda} = Var[m(\hat{\theta})]$. We also show the empirical and estimated moments in Figure 4.A.3.

Monte Carlo. We also conducted Monte Carlo experiments to increase our confidence in the estimation procedure. We simulate the choices of $N = 200$ agents with preferences θ_0 for randomly-chosen values of θ_0. For each θ_0, we start our estimation procedure at a perturbed initial value of $\theta_0 + \xi$. The minimum-distance estimator is able to back out θ_0 in our simulation experiments.

21. We use a Python implementation of this estimation routine (Gabler, 2020).

Appendix 4.C Conceptual framework

We briefly discuss choice-dated prosocial utility and conditions that imply a declining forward exchange rate. Recall that t denotes the current period, τ indexes time relative to t, $x_{t+\tau}$ denotes a dated payment to the decision-maker to be received at $t+\tau$, and $g_{t+\tau}$ represents a donation to charity that was caused at time t and will be received by the charity in τ periods. Suppose that the decision-maker's preferences are given by

$$U_t \ = \ \alpha(\boldsymbol{g}) + \sum_{\tau=0}^{\infty} D(\tau) u(x_{t+\tau}), \tag{4.C.1}$$

where $\alpha(\cdot)$ captures the choice-dated prosocial utility derived from the stream of future donations $\boldsymbol{g} = (g_{t+\tau})_\tau$ that has been *caused* in t. As we are mainly interested in the effect of delays, we replace α by a linear approximation

$$\alpha(\boldsymbol{g}) \ \approx \ a \sum_{\tau=0}^{\infty} D^c(\tau) g_{t+\tau}, \tag{4.C.2}$$

where $D^c(\tau)$ can be interpreted as an implicit "discount factor" that describes how choice-dated prosocial utility from causing a future charitable donation depreciates with the delay of the donation. We provide a sufficient condition for an asymptotically declining forward exchange rate:

Assumption 1. The implicit discount factor $D^c(\tau)$ declines at a lower rate than the subjective discount factor $D(\tau)$, i.e. $\lim_{\tau\to\infty} D^c(\tau)/D(\tau) = \infty$.

Intuitively, this implies that the choice-dated prosocial utility from the act of giving is less sensitive to the delay τ than the utility from payments to the self.[22] Thus, for large τ, the choice-dated prosocial utility will be insensitive to the delay τ *relative* to the sensitivity of utility from self-euros: the forward exchange rate will converge to zero.

We provide a simple example to illustrate why we would expect this condition to hold. Suppose that causing a delayed donation $g_{t+\tau}$ at time t provides an immediate feeling of warm glow (Andreoni, 1989), $\bar{\alpha}$, independent of the size of the donation itself, in addition to other sources of choice-dated prosocial utility, i.e. suppose that the choice-dated prosocial utility generated by $g_{t+\tau}$ is:

$$\bar{\alpha}\mathbb{1}\left(g_{t+\tau} \ > \ 0\right) + v_\tau(g_{t+\tau}), \tag{4.C.3}$$

where $v_\tau(g_{t+\tau})$ is a family of positive function. Today, the decision-maker prefers a delayed donation $g_{t+\tau}$ in τ periods to an equally delayed amount $x_{t+\tau}$ of self-euros

22. If we are willing to assume exponential discounting, i.e. $D^c(\tau) = \delta_c^\tau$ and $D(\tau) = \delta^\tau$, the assumption is equivalent to $\delta_c > \delta$.

if

$$\bar{\alpha} + v_\tau(g_{t+\tau}) \;\geq\; D(\tau)u(x) \quad \Longleftrightarrow \quad \underbrace{\frac{\bar{\alpha}}{D(\tau)u(x)}}_{\to\infty} + \underbrace{\frac{v_\tau(g_{t+\tau})}{D(\tau)u(x)}}_{\geq 0} \;\geq\; 1. \tag{4.C.4}$$

Thus, for large τ, the decision-maker will prefer the donation to contemporaneous self-euros, implying an asymptotically declining forward exchange rate. Note that we only need the existence of an (arbitrarily small) positive lower bound on the utility from the act of giving itself to obtain this result:

Proposition 6. *Suppose that the choice-dated prosocial utility from causing a dated donation g at time t that will be received by the charity at $t+\tau$ is bounded from below by $\bar{\alpha} > 0$. Then, the forward exchange rate converges to zero.*

Intuitively, the subjective discount factors imply that the present value of future self-euros becomes negligible for large τ and eventually falls below the lower bound on the immediate choice-dated prosocial utility (e.g. "warm glow"). In particular, we do not need any additional assumptions on the source of prosocial utilities.

Appendix 4.D Experimental Instructions

The original instructions used in the laboratory experiment are in German. We provide an English translation of the instructions below. The experiment has two parts. Each part consists of five different stages and each stage contains multiple price lists. To avoid repetitions, we only include the translation of one price list per stage. Within a stage, the instructions are constant across price lists except for changes in the monetary amounts or the number of months until a payment is made. See Section 4.3 of the paper for more details on how the price lists were constructed. The following sections contain the translations:

4.D.1 Introduction

Welcome and thank you for your interest in this study!

For your participation you will receive a fixed payment of 10.00 €, which will be paid to you by bank transfer after the study. In this study you will make decisions on the computer. Depending on how you decide you can earn additional money.

You are not allowed to talk to other participants during the study. Please turn off your mobile phone now, so that other participants will not be disturbed. Please only use the designated functions on the computer and make your entries using the keyboard and the mouse. If you have any questions, please raise your hand. Your question will be answered at your seat.

On the following screens you will see detailed information concerning the study. After reading this information you can confirm or refuse your participation.

To proceed click "Next".

[**end of screen**]

Information on participating in this study by the *BonnEconLab*

The following information has been sent to you via email along with the confirmation of your registration for this study. You will receive this information again now. Once you have read the subsequent declaration of consent you can confirm your participation by clicking on "I agree".

[followed by mandated exclusion restrictions for participation in this study]

[**end of screen**]

Information

In the follow part of this study, you will see important information, concerning tuberculosis and its possible treatment, that is relevant for your subsequent decisions. Please read all information carefully.

[end of screen]

Information about Tuberculosis

What is tuberculosis?

Tuberculosis – also called consumptiveness or White Death – is an infectious disease, which is caused by bacteria. Roughly one third of all humans are infected with the pathogen of tuberculosis. Active tuberculosis breaks out among 5 to 10% of all those infected. Tuberculosis is primarily airborne. This is also why a quick treatment is necessary.

What are the symptoms of tuberculosis?

Tuberculosis patients often suffer from generalized symptoms like fatigue, feeling of weakness, lack of appetite, and weight loss. At an advanced stage of lung tuberculosis, the patient coughs up blood, leading to the so-called rush of blood. Without treatment a person with tuberculosis dies with a probability of 43%.

How prevalent is tuberculosis?

In the year 2014, 6 million people have been recorded as falling ill with active tuberculosis. Almost 1.5 million people die of tuberculosis each year. This means more deaths are caused by tuberculosis than HIV, malaria, or any other infectious disease.

Is tuberculosis curable?

Today tuberculosis is curable. Treatment is administered by giving antibiotics several times each week over a period of 6 months. It is important that there is no interruption of treatment. In the years from 2000 to 2014 approximately 43 million human lives were saved due to the effective diagnosis and treatment of tuberculosis. The success rate of treatment for a new infection is often above 85%. The preceding numbers and information are provided by the World Health Organization (WHO), the United Nations' institution for the international public health, and are freely available. You can check this information on the web page of the WHO after this study.

Figure 4.D.1. Typical appearance of a tuberculosis patient

[end of screen]

Your decision

In the course of this study you can choose between options that have different consequences. In particular, you can choose between options with the following consequences:

Additional Payment: If you choose this option, you will receive an additional payment.

Saving a Human Life: If you choose this option, you will not receive an additional payment. This option has another consequence: You save one human life.

After it has emerged which option will be implemented for you, it will be carried out exactly as described. On the next tab you will receive more information about the implementation of Saving a Human Life.

[**end of screen**]

Information about saving a human life

How will a human life be saved?

Depending on how you decide, a human life can be saved. A human life will be saved by arranging a donation of 350.00 € on your behalf to an organization that identifies and treats people suffering from tuberculosis. This donation will be executed for you by the BonnEconLab after the study. The entire donation amount will be used by the organization for the direct treatment of tuberculosis.

What does it mean to "save a life"?

In this context, to save a human life means to successfully cure one person of tuberculosis, who *otherwise* would have died from the disease. This means in particular: The donation amount is sufficient to identify and cure as many sick people such that there is at least one person among them, who would otherwise have died from tuberculosis in expectation. The calculation of the amount accommodates the fact that there are other ways (e.g., the national health care system) through which people can be cured. That means: **The amount of 350.00 € was calculated in such a way that the organization can save at least one additional human from death.**

On the next tab you will receive additional information about the possible saving of a human life and details about the organization that treats tuberculosis patients.

[end of screen]

Operation ASHA

Your decisions can save a human life. Depending on how you decide, an amount of 350.00 € will be transferred to the organization *Operation ASHA* after the study.

Operation ASHA is a charity organization that has specialized in the treatment of tuberculosis in disadvantaged communities since 2005. The work of *Operation ASHA* is based on the insight that the biggest obstacle for the treatment of tuberculosis is the interruption of the necessary 6-month-long regular intake of medication. For a successful treatment the patient has to come to a medical facility twice a week – more than 60 times in total – to take the medication. An interruption or termination of the treatment is fatal, because this strongly enhances the development of a drug-resistant form of tuberculosis. This form of tuberculosis is much more difficult to treat and almost always leads to death.

To overcome this problem, *Operation ASHA* developed a concept that guarantees the regular treatment through immediate spatial proximity to the patient. A possible non-adherence is additionally prevented by visiting the patient at home. By now *Operation ASHA* runs more than 360 treatment centers, almost all of which are located in the poorest regions of India. More than 60,000 sick individuals have been identified and treated this way.

Figure 4.D.2. An employee of Operation ASHA provides medicine to a tuberculosis patient.

Operation ASHA is an internationally recognized organization, and its success has been covered by many news outlets including the New York Times, the BBC, and Deutsche Welle. MIT and University College London have already conducted research projects about the fight against tuberculosis in cooperation with *Operation ASHA*. The treatment method employed by *Operation ASHA* is described by the World Health Organization (WHO) as “highly efficient and cost-effective”.

[end of screen]

What determines the donation amount for saving a human life?

The donation amount ensures that at least one human life is saved in expectation.

The information used for the calculation of the donation amount exclusively consists of public statements by the World Health Organization (WHO), peer-reviewed research studies, statistical releases from the Indian government, and published figures from *Operation ASHA*. In the calculation all information was interpreted in a conservative way and more pessimistic estimates were used in case of doubt such that the donation amount of 350.00 € is, if anything, higher than the actual costs associated with saving a human life. Moreover, the calculation was based on the treatment success rate of *Operation ASHA* and the mortality rate of an alternative treatment by the national tuberculosis program in India. Furthermore, different detection rates for new cases of tuberculosis have been accounted for.

Based on a very high number of cases, one can illustrate the contribution of your donation as follows:

With your donation, *Operation ASHA* can treat five additional tuberculosis patients.

If these five sick individuals were not treated by *Operation ASHA*, one patient would die in expectation. If five people are treated by means of your donation, no patient dies in expectation. Based on these expected values, one human life will be saved with your donation. This relationship is depicted in the following diagram.

a) Without treatment by *Operation ASHA*, one of five individuals sick with tuberculosis will die in expectation.

b) With the donation five individuals sick with tuberculosis can be treated by *Operation ASHA*, and none of these individuals will die in expectation.

An agreement with *Operation ASHA* for the purpose of this study ensures that 100% of the donation amount will exclusively be used for the diagnosis and treatment of tuberculosis patients. That means that every euro of the donation amount will directly go toward saving human lives.

[end of screen]

Summary

Tuberculosis

The success rate of medical treatment for a new infection is very high. Nevertheless, 1.5 million people die from tuberculosis each year. The biggest obstacle for the cure of tuberculosis is a possible termination of the regular treatment with antibiotics. The concept of *Operation ASHA* is therefore based on having direct spatial proximity to its patients and being able to control and account for the regular intake of medication.

Your decision

In the course of this study you can choose between options that have different consequences. In particular, you can choose between options with the following consequences: You can choose the additional monetary payment. If you choose the other option, you will not receive an additional monetary payment, but you can save a human life. Concretely, by choosing the other option you will cause a donation. The donation of 350.00 € will be paid on your behalf, which is sufficient not only to cure one person, but to actually save that person from dying of tuberculosis.

How is the human life saved?

The donation amount of 350.00 € already accounts for the fact that a sick person could also have survived without treatment by *Operation ASHA*; or that he could instead have been treated by the national health care system. This is why the amount is sufficient for the diagnosis and complete treatment of several affected individuals.

Please note: **This is not a hypothetical game**. The option to be implemented for you will actually be carried out – exactly as described – by the *BonnEconLab*. You will receive the money in case you choose the additional monetary payment. In case you choose to save a human life, we will allow inspection of the confirmed bank transfer to the organization *Operation ASHA* upon request.

If you have individual questions, you can also direct these by email after the study to nachbesprechung@uni-bonn.de. You find this email address on the back of your seating card. You can take it home with you. Click on "Next", if you have carefully read the information on this page. Please note: You can only click on the button "Next" once you have spent at least five minutes on the seven tabs of this page.

[end of screen]

Information on the next part of this study

In the next part of this study, we will ask you to make a series of decisions in which you can choose between two monetary payments. The dates on which the two monetary payments are made can differ.

About this part of the study

This part of the study consists of five parts. In each part, you will make a decision in five different decision-making scenarios. At the beginning of each part, you will receive information that is relevant for this part. At the beginning of each decision-making scenario, you will also receive additional information for this particular decision-making scenario.

Payments in this part of the study

All monetary payments in this part of the study will be made by bank transfer. Each bank transfer will be made on the exact date that was indicated for the monetary payments. If, for example, a decision is about a monetary payment today, the corresponding monetary amount will be sent to you by a bank transfer today. If the decision involves a monetary payment in one month, a bank transfer with the corresponding amount will be made exactly one month from now.

In what follows, you will face a series of decision-making scenarios. One of these decision-making scenarios will be randomly selected by the computer at the end of this study. Your decision in *this* decision-making scenario will be implemented at the end of this study.

Remember:

- Every decision-making scenario can be relevant for your monetary payment.
- Your decisions in this part determine both **to whom** the monetary payment will go and **at which date** the monetary payment will be made.
- All monetary payments will be made by bank transfer.

[end of screen]

What does it mean that a donation will be made earlier or later?

If a donation is made earlier because of your decisions, help will be available earlier and hence people can be saved from death at an earlier point in time.

If a donation is made later, for example, in one year from now, then help will only be available later. Hence, people can only be saved from death at a later point in time. This means that the donation will be too late to help some patients that have tuberculosis in the present. In this case, patients who got sick at a later date will receive treatment instead.

The **size of the donation** is important, because more people can be helped with more money.

When making the following decisions, you should therefore take into account **when** the donation will be made and **how much** will be donated based on your decisions.

[end of screen]

4.D.2 Experiment Part A

4.D.2.1 UD-S

Information for the current part

In the following, you will see a series of decision-making scenarios in which you can choose between Option A and Option B.

- **Option A**: A **smaller** monetary payment **to you** at an **earlier** date.
- **Option B**: A **larger** monetary payment **to you** at a **later** date.

Thus, you can make a decision about a payment to yourself. You have the choice between a monetary payment that is smaller and made earlier; and a monetary payment that is larger, but made later.

Please note:

- Each of the following decisions could be the one that is actually implemented.
- **All** monetary payments will be made by **bank transfer**.

[**end of screen**]

Information for the decision-making scenario on the next page

[Box that repeats the relevant information for the current part of the study]

On the next page, you will see a list of choices between

- **Option A**: A smaller monetary payment to you today.
- **Option B**: A larger monetary payment to you in 12 months.

You can thus decide whether you are willing to wait to receive a larger monetary payment.

[**end of screen**]

You can now make your decision

Please indicate in each row of this table whether you choose **Option A** or **Option B**.

Option A	Option B
50.00 € for you today ○	○ 50.00 € for you in 12 months
50.00 € for you today ○	○ 52.50 € for you in 12 months
50.00 € for you today ○	○ 55.00 € for you in 12 months
...○	○ ...
50.00 € for you today ○	○ 120.00 € for you in 12 months
50.00 € for you today ○	○ 122.50 € for you in 12 months
50.00 € for you today ○	○ 125.00 € for you in 12 months

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[**end of screen**]

4.D.2.2 UD-C

Information for the current part

In the following, you will see a series of decision-making scenarios in which you can choose between Option A and Option B.

- **Option A**: A **smaller** monetary payment **to *Operation ASHA*** at an **earlier** date.
You are making a smaller contribution to saving lives and the contribution is made earlier.
- **Option B**: A **larger** monetary payment **to *Operation ASHA*** at a **later** date.
You are making a larger contribution to saving lives. However, the contribution is made later, so there is a delay.

Thus, you can choose whether you want to make a smaller donation at an earlier date to save fewer human lives, or whether you want to wait to make a larger donation at a later date to save more human lives.

Please note:

- Each of the following decisions could be the one that is actually implemented.
- **All** monetary payments will be made by **bank transfer**.

[**end of screen**]

Information for the decision-making scenario on the next page

[Box that repeats the relevant information for the current part of the study]

On the next page, you will see a list of choices between

- **Option A**: A smaller monetary payment to *Operation ASHA* today.
- **Option B**: A larger monetary payment to *Operation ASHA* in 12 months.

100% of the donation amount will be used to save human lives.

You can thus decide whether you prefer to save fewer human lives at an earlier date in the immediate future, or whether you want to help save more human lives in the future, but with a greater delay.

[**end of screen**]

You can now make your decision

Please indicate in each row of this table whether you choose **Option A** or **Option B**.

Option A	Option B
50.00 € for Operation ASHA today ○	○ 50.00 € for Operation ASHA in 12 months
50.00 € for Operation ASHA today ○	○ 52.50 € for Operation ASHA in 12 months
50.00 € for Operation ASHA today ○	○ 55.00 € for Operation ASHA in 12 months
...○	○ ...
50.00 € for Operation ASHA today ○	○ 120.00 € for Operation ASHA in 12 months
50.00 € for Operation ASHA today ○	○ 122.50 € for Operation ASHA in 12 months
50.00 € for Operation ASHA today ○	○ 125.00 € for Operation ASHA in 12 months

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[**end of screen**]

4.D.2.3 ER

Information for the current part

In the following, you will see a series of decision-making scenarios in which you can choose between Option A and Option B.

- **Option A**: Monetary payment **to you** at a **given** date.
- **Option B**: Monetary payment **to Operation ASHA** on the **same** date.
You are making a contribution to saving human lives on the same date that you would have received your monetary payment if you had chosen **Option A**.

Thus, you can choose whether you prefer making a monetary payment to yourself on a given date, or whether you prefer making a donation to help save human lives on the same date.

Please note:

- Each of the following decisions could be the one that is actually implemented.
- **All** monetary payments will be made by **bank transfer**.

[**end of screen**]

Information for the decision-making scenario on the next page

[Box that repeats the relevant information for the current part of the study]

On the next page, you will see a list of choices between

- **Option A**: A monetary payment to *you* in 12 months.
- **Option B**: A monetary payment to *Operation ASHA* in 12 months.

100% of the donation amount will be used to save human lives.

You can thus decide whether you are willing to forego a monetary payment to yourself in 12 months in order to save human lives.

[**end of screen**]

You can now make your decision

Please indicate in each row of this table whether you choose **Option A** or **Option B**.

Option A	Option B
50.00 € for you in 12 months ○	○ 0.00 € for Operation ASHA in 12 months
50.00 € for you in 12 months ○	○ 10.00 € for Operation ASHA in 12 months
50.00 € for you in 12 months ○	○ 20.00 € for Operation ASHA in 12 months
…○	○ …
50.00 € for you in 12 months ○	○ 180.00 € for Operation ASHA in 12 months
50.00 € for you in 12 months ○	○ 190.00 € for Operation ASHA in 12 months
50.00 € for you in 12 months ○	○ 200.00 € for Operation ASHA in 12 months

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[**end of screen**]

4.D.2.4 MD-S

Information for the current part

In the following, you will see a series of decision-making scenarios in which you can choose between Option A and Option B.

- **Option A**: A monetary payment **to you** at an **earlier** date.
- **Option B**: A monetary payment **to Operation ASHA** at a **later** date.

You are making a contribution to saving lives. However, the contribution is made later, so there is a delay.

Thus, you can choose whether you prefer a monetary payment to yourself at an earlier date, or whether you prefer to wait to make a larger donation to help save human lives at a later date.

Please note:

- Each of the following decisions could be the one that is actually implemented.
- **All** monetary payments will be made by **bank transfer**.

[**end of screen**]

Information for the decision-making scenario on the next page

[Box that repeats the relevant information for the current part of the study]

On the next page, you will see a list of choices between

- **Option A**: A monetary payment to *you* today.
- **Option B**: A monetary payment to *Operation ASHA* in 12 months.

100% of the donation amount will be used to save human lives.

You can thus decide whether you are willing to forego a monetary payment to yourself at an earlier date to save human lives at a later date.

[**end of screen**]

You can now make your decision

Please indicate in each row of this table whether you choose **Option A** or **Option B**.

Option A	Option B
50.00 € for you today ○	○ 0.00 € for Operation ASHA in 12 months
50.00 € for you today ○	○ 15.00 € for Operation ASHA in 12 months
50.00 € for you today ○	○ 30.00 € for Operation ASHA in 12 months
...○	○ ...
50.00 € for you today ○	○ 345.00 € for Operation ASHA in 12 months
50.00 € for you today ○	○ 360.00 € for Operation ASHA in 12 months
50.00 € for you today ○	○ 375.00 € for Operation ASHA in 12 months

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[**end of screen**]

4.D.2.5 MD-C

Information for the current part

In the following, you will see a series of decision-making scenarios in which you can choose between Option A and Option B.

- **Option A**: A monetary payment **to Operation ASHA** at an **earlier** date.
  You are making a contribution to saving lives at an earlier date.
- **Option B**: A monetary payment **to you** at a **later** date.

Thus, you can choose whether you prefer a donation to help save human lives at an earlier date, or whether you prefer to wait to receive a monetary payment for yourself at a later date.

Please note:

- Each of the following decisions could be the one that is actually implemented.
- **All** monetary payments will be made by **bank transfer**.

[**end of screen**]

Information for the decision-making scenario on the next page

[Box that repeats the relevant information for the current part of the study]

On the next page, you will see a list of choices between

- **Option A**: A monetary payment to *Operation ASHA* today.
- **Option B**: A monetary payment to *you* in 12 months.

100% of the donation amount will be used to save human lives.

You can thus decide whether you are willing to forego saving human lives at an earlier date to receive a monetary payment at a later date.

[**end of screen**]

You can now make your decision

Please indicate in each row of this table whether you choose **Option A** or **Option B**.

Option A	Option B
50.00 € for Operation ASHA today ○	○ 0.00 € for you in 12 months
50.00 € for Operation ASHA today ○	○ 5.00 € for you in 12 months
50.00 € for Operation ASHA today ○	○ 10.00 € for you in 12 months
...○	○ ...
50.00 € for Operation ASHA today ○	○ 115.00 € for you in 12 months
50.00 € for Operation ASHA today ○	○ 120.00 € for you in 12 months
50.00 € for Operation ASHA today ○	○ 125.00 € for you in 12 months

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[**end of screen**]

4.D.3 Experiment Part B

Task description

In the following part of the study, we ask you make a series of decisions involving a choice between two lotteries, **Lottery A** and **Lottery B**. Both lotteries will be determined by a fair coin toss. That means that there is a 50% chance that it lands on heads, and a 50% chance that it lands on tails.

Before each lottery choice, you will receive information about the initial endowment in this decision. This initial endowment consists of two parts:

- A monetary payment **to you**
- A monetary payment **to Operation ASHA**. 100% of this amount will be used to save human lives.

After you have received information about the initial endowment, you can make your choice between **Lottery A** and **Lottery B**.

Please note:

- The lotteries will change the monetary payments to you and/or the organization. You will learn exactly how the initial endowments will change if, for example, you choose Lottery A and the coin toss lands on heads.
- **Thus, how the monetary payments to you and the organization change depends both on which lottery you choose and the result of the coin toss.** The coin toss will be carried out by the computer.

Payments in this part of the study

All monetary payments in this part of the study will be made by bank transfer. In the following decision-making scenarios, monetary payments are made either to you or to the organization *Operation ASHA*. If you are the recipient, a bank transfer to your account will be made today. If *Operation ASHA* is the recipient of the monetary payment, a bank transfer to the organization's account will be made today. As previously explained, 100% of the amount that is transferred to the organization's account will be used to save people from dying of tuberculosis.

In what follows, you will face a series of decision-making scenarios. One of these decision-making scenarios will be randomly selected by the computer at the end of this study. Your decision in *this* decision-making scenario will be implemented by a bank transfer at the end of this study. Your decisions in this part of the study thus determine *which* lottery is played at the end of this study.

Remember:

- Every decision-making scenario can be relevant for your monetary payment.

- Your decisions in this part determine both **to whom** the monetary payment will go and **at which date** the monetary payment will be made.

- All monetary payments will be made by bank transfer.

[end of screen]

Example

In the following decision-making scenarios, you can choose between **Lottery A** and **Lottery B**. On this page, we use an example to illustrate the choice between both lotteries.

In the following decision-making scenarios, you will see a page that looks like this:

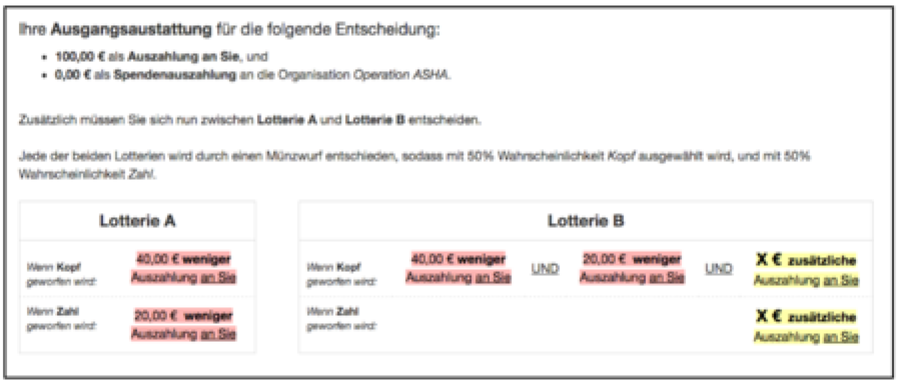

On such a page, you will see information about the initial endowment, and how these endowments change depending on which lottery you choose and what the result of the coin toss is.

In the picture below, we explain the elements of this page in more detail:

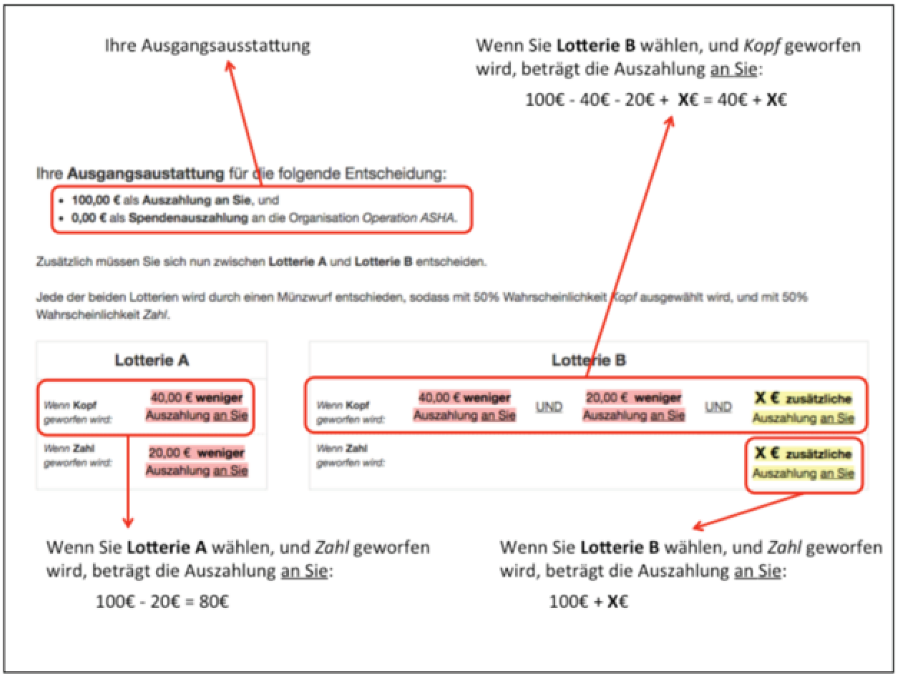

In each decision-making scenario where you have to choose between **Lottery A** and **Lottery B**, we will show you an amount **X €**. The picture below illustrates what your decision would look like if $X = 10.00$ €. By selecting the left or right circle, you can choose between **Lottery A** and **Lottery B**.

To proceed click "Next".

[end of screen]

Exercise 1

On this and the following page, you can check whether you have correctly understood all the necessary information for this part of the study. For the first exercise, take a look at the following initial endowment:

The initial endowment for the following scenario:

- *25.00 € for you, and*

- *a donation of 25.00 € to the organization Operation ASHA.*

*In addition, you also have to choose between **Lottery A** and **Lottery B**.*

Imagine that, given the initial endowment above, you had to make a decision between the following two lotteries:

- Lottery A:
  - If the coin toss is heads: the donation amount is reduced by 10.00 €.
  - If the coin toss is tails: the monetary payment to you is reduced by 10.00 €.
- Lottery B:
  - If the coin toss is heads: both the donation amount and the monetary payment to you are reduced by 10.00 €. You receive an additional **X** € as well.
  - If the coin toss is tails: you receive an additional **X** €.
  - $X = 2.00$ €

To test whether you have understood how your choice between Lottery A and Lottery B as well as how the outcome of the coin toss affects the monetary payments, please provide answers to the following questions:

- If I choose **Lottery A** and the coin toss is *heads*, the monetary amount that I will receive, including the initial endowment, is: [blank field] (in €)
- If I choose **Lottery B** and the coin toss is *heads*, the monetary amount that I will receive, including the initial endowment, is: [blank field] (in €)
- If I choose **Lottery B** and the coin toss is *heads*, the size of the donation, including the initial endowment, is: [blank field] (in €)
- If I choose **Lottery B** and the coin toss is *tails*, the monetary amount that I will receive, including the initial endowment, is: [blank field] (in €)

[end of screen]

Exercise 2

For the first exercise, take a look at the following initial endowment:

The initial endowment for the following scenario:

- *40.00 € for you, and*
- *a donation of 0.00 € to the organization Operation ASHA.*

In addition, you also have to choose between ***Lottery A*** *and* ***Lottery B****.*

Some decisions involve a so-called **additional lottery**. Every additional lottery has a possible positive outcome (the monetary payment increases) and a possible negative outcome (the monetary payment decreases). The outcome of the **additional lottery** will also be randomly determined by the computer.

Note: Pay attention to the probabilities in the **additional lottery**.

Imagine that, given the initial endowment above, you had to make a decision between the following two lotteries:

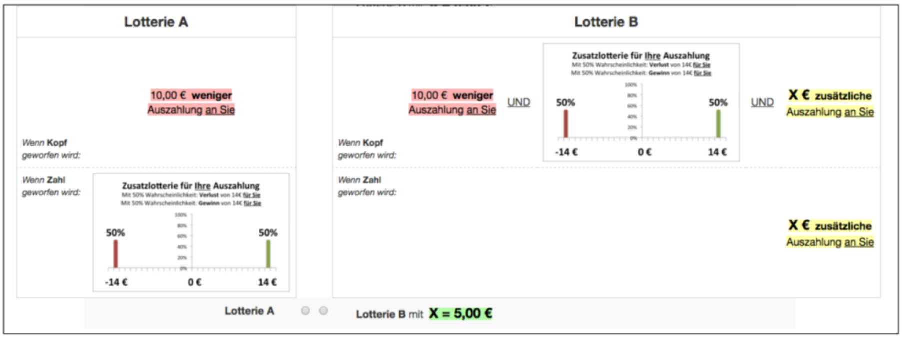

- Lottery A:
  - If the coin toss is heads: the donation amount is reduced by 10.00 €.
  - If the coin toss is tails: There is an additional lottery for **your** monetary payment.
    - With a probability of 50%: You lose 14 €.
    - With a probability of 50%: You win 14 €.
- Lottery B:
  - If the coin toss is heads: the donation amount is reduced by 10.00 € AND you will receive an additional **X** € AND have an additional lottery for **your** monetary payment:
    - With a probability of 50%: You lose 14 €.
    - With a probability of 50%: You win 14 €.
  - If the coin toss is tails: you receive an additional **X** €.
  - $X = 5.00$ €

The additional lottery thus has a possible negative outcome of -14.00 € and a possible positive outcome of +14.00 €. Both outcomes are equally likely, that is, they both have a probability of 50%.

To test whether you have understood how your choice between Lottery A and Lottery B as well as how the outcome of the coin toss affects the monetary payments, please provide answers to the following questions:

- If I choose **Lottery A** and the coin toss is tails, then the outcome of the additional lottery is +14 €, and I will receive a monetary payment, including the initial endowment, of: [blank field] (in €)
- If I choose **Lottery B** and the coin toss is heads, then the outcome of the additional lottery is -14 €, and I will receive a monetary payment, including the initial endowment, of: [blank field] (in €)

[end of screen]

Your task begins on the next page

On the next page you will see the first decision-making scenario. From now on, the decisions you make are no longer an exercise, meaning that any of your following decisions and all related consequences could be implemented.

Remember:

- Every decision-making scenario can be relevant for your monetary payment.
- Your decisions in this part determine both **to whom** the monetary payment will go and **at which date** the monetary payment will be made.
- All monetary payments will be made by bank transfer.

To proceed click "Next".

4.D.3.1 RA–Self

The initial endowment for this decision is:

- 25.00 € for you, and
- a donation of 0.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

Both lotteries will be decided by a coin toss, which means that there is a 50% chance of heads and a 50% chance of tails.

[Description of the lotteries]

On the next page you will see a list where each row represents a different decision-making scenario between Lottery A and Lottery B. Each row indicates the value of **X** in that particular decision-making scenario. To proceed click "Next".

[end of screen]

Decision

The initial endowment for this decision is:

- 25.00 € for you, and
- a donation of 0.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

[Description of the lotteries]

Note: **X** € will be paid to you whenever you choose **Lottery B**, independently of whether the coin toss is heads or tails. Whether **X** is positive (a gain) or negative (a loss) depends on the decision-making scenario.

Lottery A ○	○ **Lottery B** with **X = -5.00 €**
Lottery A ○	○ **Lottery B** with **X = -4.50 €**
Lottery A ○	○ **Lottery B** with **X = -4.00 €**
...○	○ ...
Lottery A ○	○ **Lottery B** with **X = 4.00 €**
Lottery A ○	○ **Lottery B** with **X = 4.50 €**
Lottery A ○	○ **Lottery B** with **X = 5.00 €**

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[end of screen]

4.D.3.2 RA–Charity

The initial endowment for this decision is:

- 0.00 € for you, and
- a donation of 25.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

Both lotteries will be decided by a coin toss, which means that there is a 50% chance of heads and a 50% chance of tails.

[Description of the lotteries]

On the next page you will see a list where each row represents a different decision-making scenario between Lottery A and Lottery B. Each row indicates the value of **X** in that particular decision-making scenario. To proceed click "Next".

[end of screen]

Decision

The initial endowment for this decision is:

- 0.00 € for you, and
- a donation of 25.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

[Description of the lotteries]

Note: **X** € will be paid to you whenever you choose **Lottery B**, independently of whether the coin toss is heads or tails. Whether **X** is positive (a gain) or negative (a loss) depends on the decision-making scenario.

Lottery A ○	○ **Lottery B** with **X = -5.00 €**
Lottery A ○	○ **Lottery B** with **X = -4.50 €**
Lottery A ○	○ **Lottery B** with **X = -4.00 €**
...○	○ ...
Lottery A ○	○ **Lottery B** with **X = 4.00 €**
Lottery A ○	○ **Lottery B** with **X = 4.50 €**
Lottery A ○	○ **Lottery B** with **X = 5.00 €**

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[end of screen]

4.D.3.3 X–RA

The initial endowment for this decision is:

- 25.00 € for you, and
- a donation of 25.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

Both lotteries will be decided by a coin toss, which means that there is a 50% chance of heads and a 50% chance of tails.

[Description of the lotteries]

On the next page you will see a list where each row represents a different decision-making scenario between Lottery A and Lottery B. Each row indicates the value of **X** in that particular decision-making scenario. To proceed click "Next".

[end of screen]

Decision

The initial endowment for this decision is:

- 25.00 € for you, and
- a donation of 25.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

[Description of the lotteries]

Note: **X** € will be paid to you whenever you choose **Lottery B**, independently of whether the coin toss is heads or tails. Whether **X** is positive (a gain) or negative (a loss) depends on the decision-making scenario.

Lottery A ○	○ **Lottery B** with **X = -5.00 €**
Lottery A ○	○ **Lottery B** with **X = -4.50 €**
Lottery A ○	○ **Lottery B** with **X = -4.00 €**
...○	○ ...
Lottery A ○	○ **Lottery B** with **X = 4.00 €**
Lottery A ○	○ **Lottery B** with **X = 4.50 €**
Lottery A ○	○ **Lottery B** with **X = 5.00 €**

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.

[end of screen]

4.D.3.4 PR–Self

The initial endowment for this decision is:

- 40.00 € for you, and
- a donation of 0.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

Both lotteries will be decided by a coin toss, which means that there is a 50% chance of heads and a 50% chance of tails.

[Description of the lotteries]

This decision entails the possibility of an **additional lottery**. For example, if you choose Lottery A and the coin toss is tails, the additional lottery will be played. The outcome of the additional lottery will be determined by the computer.

On the next page you will see a list where each row represents a different decision-making scenario between Lottery A and Lottery B. Each row indicates the value of **X** in that particular decision-making scenario. To proceed click "Next".

[end of screen]

Decision

The initial endowment for this decision is:

- 40.00 € for you, and
- a donation of 0.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

[Description of the lotteries]

Note: **X** € will be paid to you whenever you choose **Lottery B**, independently of whether the coin toss is heads or tails. Whether **X** is positive (a gain) or negative (a loss) depends on the decision-making scenario.

Lottery A ○	○ **Lottery B** with **X = -5.00 €**
Lottery A ○	○ **Lottery B** with **X = -4.50 €**
Lottery A ○	○ **Lottery B** with **X = -4.00 €**
...○	○ ...
Lottery A ○	○ **Lottery B** with **X = 4.00 €**
Lottery A ○	○ **Lottery B** with **X = 4.50 €**
Lottery A ○	○ **Lottery B** with **X = 5.00 €**

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don’t have to click as much.

[end of screen]

4.D.3.5 PR–Charity

The initial endowment for this decision is:

- 0.00 € for you, and
- a donation of 40.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

Both lotteries will be decided by a coin toss, which means that there is a 50% chance of heads and a 50% chance of tails.

[Description of the lotteries]

This decision entails the possibility of an **additional lottery**. For example, if you choose Lottery A and the coin toss is tails, the additional lottery will be played. The outcome of the additional lottery will be determined by the computer.

On the next page you will see a list where each row represents a different decision-making scenario between Lottery A and Lottery B. Each row indicates the value of **X** in that particular decision-making scenario. To proceed click "Next".

[end of screen]

Decision

The initial endowment for this decision is:

- 0.00 € for you, and
- a donation of 40.00 € to the organization Operation ASHA.

In addition, you also have to choose between **Lottery A** and **Lottery B**.

[Description of the lotteries]

Note: **X** € will be paid to you whenever you choose **Lottery B**, independently of whether the coin toss is heads or tails. Whether **X** is positive (a gain) or negative (a loss) depends on the decision-making scenario.

Lottery A ○	○ **Lottery B** with **X = -5.00 €**
Lottery A ○	○ **Lottery B** with **X = -4.50 €**
Lottery A ○	○ **Lottery B** with **X = -4.00 €**
...○	○ ...
Lottery A ○	○ **Lottery B** with **X = 4.00 €**
Lottery A ○	○ **Lottery B** with **X = 4.50 €**
Lottery A ○	○ **Lottery B** with **X = 5.00 €**

Automatic completion: We have activated a fill-in aid that automatically fills out the remaining rows so you don't have to click as much.