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

Christian Rudolf Apenbrink

aus Düsseldorf

2023

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Florian Zimmermann
Zweitreferent:	Prof. Dr. Lorenz Goette
Tag der mündlichen Prüfung:	15. März 2023

Acknowledgements

During the process of writing this dissertation, I have benefited tremendously from the support of both the academic community in Bonn and my social circle.

First of all, I am deeply indebted to Florian, who became my advisor just after joining the University of Bonn as an assistant professor. Throughout the years, he has been a constant source of inspiration, encouragement, and support for me, and I admire his dedication to being available for his graduate students. Likewise, I am extremely grateful to Lorenz, my secondary advisor, for his continued guidance and the positivity he displayed during our meetings and discussions, which always left me feeling confident about taking the next steps.

The Bonn Graduate School of Economics, German Academic Exchange Service, Institute for Applied Microeconomics, and Institute on Behavior and Inequality provided financial support, without which this dissertation might not exist. I also want to thank the staff at these institutions, especially Andrea and Simone, for their assistance with various administrative matters.

Rising to the academic and personal challenges of a PhD was eased by the company of various colleagues and friends at the Bonn Graduate School of Economics, of which Chui, Laurenz, Matthias, Paul, Thomas, Tobi, and Ximi are just a few. I am especially thankful that I had Franzi and Luca by my side during an exciting first PhD year in Berkeley, which laid the foundation for my own research projects. Moreover, I want to express my gratitude to Timo for being my coauthor on a great project which eventually failed. I learned a lot from sharing this experience with him.

Finally, I am profoundly grateful to my parents and my sister for their unconditional love and support. I dedicate this dissertation to them.

Bonn, June 2022

Contents

List of Figures ix

List of Tables xii

Introduction 1

References 4

1 Experienced Political Competition and Voter Apathy 5

1.1 Introduction 5

1.2 Hypothesis in the Context of the US 10

1.3 Data 11

1.3.1 Construction of Political Competition Experiences 12

1.3.2 Sample 13

1.4 Empirical Strategy 15

1.5 Results 17

1.5.1 Preliminary Evidence 17

1.5.2 Main Results 18

1.5.3 Causal Interpretation of the Main Results 20

1.5.4 Robustness 21

1.5.5 Heterogeneity 22

1.5.6 Additional Results 24

1.6 Mechanisms 26

1.6.1 Is the Effect Driven by Experienced Political Participation? 26

1.6.2 Parental Voting Behavior 27

1.6.3 Political Identity 28

1.6.4 Civic Duty 29

1.6.5 Habit Formation 30

1.6.6 Beliefs about the Benefit of Voting 31

1.7 Conclusion 33

Appendix 1.A Data Appendix 35

1.A.1 Demographic Control Variables 35

1.A.2 Clustering to Account for Sampling Design 35

1.A.3 Historical Control Variables 36
1.A.4 Other Outcomes 37
Appendix 1.B Additional Tables 38
Appendix 1.C Robustness 43
References 49

2 The Cost of Worrying about an Epidemic: Ebola Concern and Cognitive Function in the US 57
2.1 Introduction 57
2.2 Background 62
2.3 Data 64
2.3.1 Data on Cognitive Function 64
2.3.2 Measuring Ebola Concern 66
2.3.3 Other Data Sources 68
2.3.4 Main Sample 68
2.4 OLS Estimation 68
2.5 IV Estimation 71
2.5.1 Ebola Concern and Distance to US Ebola Locations 72
2.5.2 Falsification Tests for the Exclusion Restriction 76
2.5.3 IV Estimates 76
2.5.4 Robustness 79
2.6 Do Interview Characteristics Respond to US Ebola Cases? 80
2.7 Discussion 85
2.8 Conclusion 86
Appendix 2.A HRS Data Appendix 88
Appendix 2.B Additional Tables 90
Appendix 2.C Robustness 92
References 98

3 The Cost of Worrying about an Epidemic: Experimental Evidence on Labor Productivity during the COVID-19 Pandemic 105
3.1 Introduction 105
3.2 Experimental Design 110
3.2.1 Cognitive Task 111
3.2.2 Main Experimental Conditions: *Worry-Amplifying* and *Worry-Alleviating* 113
3.2.3 Secondary Experimental Conditions: *COVID-19 Headlines* and *Neutral Headlines* 114
3.2.4 Elicitation of Additional Outcomes 115
3.2.5 Procedures 116

3.3 Preregistered Hypotheses and Analyses 119
3.3.1 Primary Hypotheses 119
3.3.2 Other Preregistered Analyses 121
3.4 Results 122
3.4.1 Baseline Sample Characteristics 122
3.4.2 Manipulation Check: Emotional Responses to the Main Treatment Manipulation 123
3.4.3 Effect of the Main Treatment Manipulation on Labor Productivity 124
3.4.4 News Salience and Labor Productivity 130
3.4.5 Evidence for Worry-Induced Attentional Capture 131
3.4.6 Effect of the Main Treatment Manipulation on Self-Perceived Productivity 133
3.5 Discussion 135
3.5.1 Are the Results Affected by Cheating on the Cognitive Task? 135
3.5.2 Differences to Findings from Previous Research 136
3.6 Conclusion 139
Appendix 3.A Additional Figures and Tables 141
Appendix 3.B Are Emotional Responses to the Videos Driven by Changes in Beliefs? 152
Appendix 3.C Robustness 158
Appendix 3.D Experimental Instructions 160
3.D.1 Instructions for Session 1 160
3.D.2 Instructions for Session 2 162
Appendix 3.E Screenshots of the Experimental Interface 169
References 170

List of Figures

1.1	Experienced Political Competition across Election Cohorts and States	14
1.2	Residual Turnout by Experienced Political Competition	18
1.C.1	Effect of Low Experienced Political Competition with Historical Control Variables	44
2.1	Ebola Case Locations and Geographic Variation in Ebola Concern across Media Markets by Week	73
3.1	Chronology of the Main Components of the Experimental Design by Experimental Condition	118
3.2	Labor Productivity in Non-news Blocks by Experimental Condition	126
3.3	Heterogeneity in the Main Treatment Effect on Labor Productivity	128
3.A.1	Evolution of the Main Treatment Effect over Time	141
3.A.2	Attention in Non-news Blocks by Experimental Condition	142
3.A.3	Working Memory in Non-news Blocks by Experimental Condition	143
3.A.4	Heterogeneity in the Effect of the Main Treatment Manipulation on Worry and Happiness	144
3.A.5	Perceived Productivity in Non-news Blocks and News Blocks by Experimental Condition	145
3.B.1	Beliefs about COVID-19 Infection Risk by Experimental Condition	154
3.B.2	Beliefs about the Risk of Long-Term Effects after a COVID-19 Infection by Experimental Condition	155
3.B.3	Beliefs about the Risk of Death Due to a COVID-19 Infection by Experimental Condition	156
3.B.4	Beliefs about the Number of Months until COVID-19 Vaccination by Experimental Condition	157
3.E.1	Sequence of Screens for Each Mental Arithmetic Problem of the Cognitive Task	169

List of Tables

1.1	Main Results	19
1.2	Heterogeneity	24
1.3	Experienced Political Competition and Experienced Turnout	27
1.4	Evidence on Mechanisms: Parental Voting Behavior, Political Identity, and Civic Duty	29
1.5	Evidence on Mechanisms: Beliefs about the Benefit of Voting	33
1.B.1	Descriptive Statistics for the Main Sample	38
1.B.2	Subsample Analyses for Turnout in Presidential and Midterm Elections	39
1.B.3	Subsample Analyses for Non-movers and Movers	40
1.B.4	Experiences in Different Periods of Life	41
1.B.5	Other Outcomes	42
1.C.1	Alternative Fixed Effects Specifications	43
1.C.2	Alternative Measures of Experienced Political Competition	45
1.C.3	Alternative Clustering of Standard Errors	46
1.C.4	Other Robustness Checks	47
1.C.5	Experiences in Different Periods of Life (Non-movers Only)	48
2.1	Ebola Concern and Cognitive Function: OLS Estimates	69
2.2	Ebola Concern and Distance to US Ebola Locations: First-Stage Estimates	74
2.3	Distance to US Ebola Locations and Selected Outcomes: Falsification Tests	77
2.4	Ebola Concern and Cognitive Function: IV Estimates	78
2.5	Distance to US Ebola Locations and Interview Characteristics	82
2.B.1	Descriptive Statistics for the Main Sample	90
2.B.2	Distance to US Ebola Locations and Interview Characteristics: Placebo Regressions	91
2.C.1	Alternative Measures of Ebola Concern: Search Interest in the Interview Week	92
2.C.2	Alternative Measures of Ebola Concern: Search Interest in the Last Two Weeks	93
2.C.3	Alternative Aggregation of Google Trends Data	94
2.C.4	Estimates for the Subsample of Metropolitan Area Interviews	95
2.C.5	Alternative Instruments: Distance to the First US Ebola Location	96

2.C.6 Alternative Instruments: Distance to the Closest of the First Two US Ebola Locations 97

3.1 Manipulation Check: Effect of the Main Treatment Manipulation on Worry and Happiness 123
3.2 Effect of the Main Treatment Manipulation on Labor Productivity in Non-news Blocks 125
3.3 Changes in Worry and Happiness and Changes in Labor Productivity 127
3.4 Effect of the Main Treatment Manipulation on Selected Outcomes 130
3.5 Effect of the Main Treatment Manipulation on News Interest 132
3.6 Effect of the Main Treatment Manipulation on Measures of Self-Perceived Productivity 134
3.A.1 Descriptive Statistics for the Analysis Sample 146
3.A.2 Descriptive Statistics by Experimental Condition and Test of Balance 147
3.A.3 Effect of the Main Treatment Manipulation on Measures of Attention 148
3.A.4 Effect of the Main Treatment Manipulation on Measures of Working Memory 149
3.A.5 Matrix of Correlation Coefficients between Changes in Labor Productivity, Worry, and Happiness 149
3.A.6 News Salience and Labor Productivity 150
3.A.7 News Salience and Labor Productivity (without Controlling for Baseline Worry) 151
3.B.1 Effect of the Main Treatment Manipulation on Beliefs about the Pandemic 153
3.B.2 Effect of the Main Treatment Manipulation on Beliefs about the Pandemic (II) 153
3.C.1 Effect of the Main Treatment Manipulation on Labor Productivity in Non-news Blocks (Alternative Control Variables) 158
3.C.2 Main Results (excluding Suspected Cheaters) 159

Introduction

Behavioral economics has enriched our understanding of economic behavior by incorporating insights from psychology and related disciplines into the standard economic model. This inflow of ideas has led to advances in methodology, like the use of survey measures and controlled experiments, as well as the development of novel concepts like limited self-control (e.g., Laibson, 1997), reference dependence (Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006), social preferences (e.g., Charness and Rabin, 2006), and systematic biases in belief formation (e.g., Tversky and Kahneman, 1974), which conform more closely to empirical observation. At the same time, the scope of economic research has expanded as behavioral economic tools and ideas have been applied to questions in a variety of domains. Often, the application of behavioral economic concepts complements standard economic analysis by uncovering new facets of or providing different perspectives on observed phenomena. Prominent examples include the analysis of labor supply (e.g., Camerer, Babcock, Loewenstein, and Thaler, 1997; Fehr and Goette, 2007) and consumption-savings decisions (e.g., Laibson, 1997). In terms of approach, a key focus of behavioral economics is on thorough empirical testing. This lays the foundation for the construction of more realistic models of behavior, with the ultimate goal of informing public policy.

This thesis follows the behavioral economic tradition by combining economic analysis, ideas from psychology, and rigorous empirical methodology to provide new perspectives on two important societal and economic issues: the determinants of voter apathy and productivity during an epidemic. It consists of three self-contained chapters, which are summarized below. While their topical focus is different, they all share the same scientific approach: For a given application, I first identify relevant psychological forces that generate novel causal hypotheses. In a second step, I then test these hypotheses empirically by developing a suitable identification strategy. As a result, my investigations make use of a variety of empirical methods that span the whole range of the toolkit, from a controlled experiment to instrumental variables estimation.

Chapter 1 addresses the determinants of observed variation in individuals' propensity to vote in elections. Understanding the causes of voter apathy is important because political participation influences election outcomes and public policy, and low or unequal turnout undermines the democratic ideal of equal representation. The starting point for my investigation is a line of research in social psychology: the so-called impressionable years hypothesis suggests that political attitudes are shaped by personal experiences made during early adulthood (e.g., Krosnick and Alwin, 1989). Building on this insight, I examine whether the relevance of early-adulthood experiences extends to the act of voting itself. I focus on a key comparative static of many economic models of voting: the competitiveness of the election (e.g., Downs, 1957). In particular, I test the hypothesis that the degree of political competition individuals experience during their impressionable years affects their propensity to vote later in life.

My identification strategy relies on cross-sectional and time-series variation in political competition in US national elections. In combination with data from a large repeated cross-sectional survey, this variation allows me to use a fixed effects approach to control for unobservable confounding factors related to the current electoral context, birth cohorts, age, and the state individuals grew up in. Holding all else equal, I then compare the self-reported turnout decisions of individuals who grew up in more or less competitive conditions, as measured by whether different or the same party won their home state in the two presidential elections experienced during the impressionable years.

I find that on average, individuals who grew up during periods of low political competition in their state are 1.4 percentage points less likely to vote later in life. This result is robust to the inclusion of a wide range of time-varying historical state characteristics, thus supporting a causal interpretation. The effect is stronger for individuals with lower education or lower income and diminishes over time, enduring for more than 20 years. These findings imply that periods of reduced political competition can have lasting negative consequences not only through their direct impact on implemented government policy (e.g., Besley, Persson, and Sturm, 2010), but also by undermining and distorting future political participation.

Within the limits of the available data, I also explore potential underlying mechanisms of the documented effect. I do not find any evidence for an intergenerational transmission of parental voting behavior or an increase in the strength of individuals' political identity. Rather, the experienced degree of electoral competition seems to persistently affect beliefs about the value of voting.

Chapter 2 and Chapter 3 are concerned with the consequences of negative emotional responses to epidemics. Existing research in economics highlights that fear of contagion can stifle economic activity by triggering avoidance behavior (e.g., Goolsbee and Syverson, 2021). However, cognitive psychology also suggests another causal pathway by which negative emotions could affect economic outcomes. In particular, there is evidence that negative emotions like worry and fear can shift scarce attentional resources towards the perceived threat, thus reducing the available cognitive capacity for threat-unrelated processing tasks (e.g., Sari, Koster, and Derakshan, 2017). Motivated by this insight, I develop the hypothesis that there is an economic cost to epidemic-induced worrying in the form of suboptimal economic choices and reduced labor productivity.

In Chapter 2, I exploit natural variation in the spread of an infectious disease to test this hypothesis with respect to cognitive function itself. An impairment of cognitive functioning is a necessary condition for a worry-induced deterioration of economic outcomes. The background of my investigation is an episode of heightened public concern about Ebola in the US in October 2014. I use data from cognitive tests administered as part of a wave of survey interviews by a large US panel study, which I combine with measures of local concern about Ebola based on internet search volume. For identification, I take advantage of temporal and spatial variation in Ebola concern caused by the emergence of four cases of Ebola in the US. Using proximity to the US cases as an instrumental variable, I show that the local level of Ebola concern individuals are exposed to at the time and place of the survey interview reduces their scores on the cognitive test. Comparing the level of Ebola concern in October 2014 with the counterfactual situation in which public concern is equal to the pre-October search interest peak, the extrapolated effect of worrying would amount to a 0.4-standard-deviation drop in test performance.

Additional analyses suggest that these results represent a direct causal effect of Ebola concern on cognitive function. In particular, I find no indication of fear-induced selection into taking the survey that could plausibly explain these results. Moreover, proximity to subsequent Ebola locations is unrelated to test scores for interviews conducted before the emergence of the first US case, which lends credence to the exclusion restriction. Taken together, my findings indicate that emotional responses to epidemics can entail a temporary cognitive cost even for individuals for whom the actual health risk never materializes. Since cognitive function is an important determinant of economic decision-making and productivity, this points to the possibility that worrying about an epidemic has adverse economic consequences.

Chapter 3 tests whether negative emotional responses to the COVID-19 pandemic impair labor productivity. This complements the previous chapter by investigating the effect of worry in a different context and focusing on a specific economic outcome that might be affected by epidemic-induced cognitive effects. A major difference lies in the chosen identification strategy. While the previous chapter exploits natural variation in emotions, this one relies on a controlled manipulation implemented as part of an online experiment. In the experiment, subjects can work on a cognitively demanding mental arithmetic task for a piece-rate wage. Before they start working, they are exposed to real media reports selected to either amplify or alleviate their worries about COVID-19.

In line with the conjecture that the COVID-19 pandemic is a major source of emotional distress, the treatment manipulation induces a multi-faceted negative emotional response, comprising an increase in worry and a decrease in happiness. However, the induced emotional responses do not translate into economically meaningful differences in labor productivity. Nonetheless, I find suggestive evidence of changes in cognition. Consistent with worry-induced attentional capture, exposure to the worry-amplifying media report increases relative interest in pandemic-related news and the reported incidence of distracting thoughts during the task. One plausible interpretation for the combined set of results is that subjects compensate for worry-induced cognitive effects by increasing their mental effort, in line with the notion of income targeting.

These findings qualify the economic significance of cognitive effects of epidemic-induced emotions, implying that they either require levels of worry that are beyond the scope of an experimental manipulation or do not necessarily translate into reductions in labor productivity. They are also relevant for public policy because they indicate that exposure to information about the danger of COVID-19 does not have adverse side effects in terms of reduced economic productivity.

Taken as a whole, this thesis produces insights of two kinds. First, it enriches our understanding of the general relevance of personal experiences and emotions for economic outcomes. Second, it has implications on the level of the two investigated issues, offering new perspectives on the determinants of voter apathy and the economic consequences of epidemics. Yet, the dissimilar results of the last two chapters illustrate that the significance of each individual scientific study is limited. There is still a lot to be learned.

References

Besley, Timothy, Torsten Persson, and Daniel M. Sturm. 2010. “Political Competition, Policy and Growth: Theory and Evidence from the US.” *Review of Economic Studies* 77 (4): 1329–52. DOI: 10.1111/j.1467-937X.2010.00606.x. [2]

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. “Labor Supply of New York City Cabdrivers: One Day at a Time.” *Quarterly Journal of Economics* 112 (2): 407–41. DOI: 10.1162/003355397555244. [1]

Charness, Gary, and Matthew Rabin. 2006. “Understanding Social Preferences with Simple Tests.” *Quarterly Journal of Economics* 117 (3): 817–69. DOI: 10.1162/003355302760193904. [1]

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York, NY: Harper & Row. [1]

Fehr, Ernst, and Lorenz Goette. 2007. “Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment.” *American Economic Review* 97 (1): 298–317. DOI: 10.1257/aer.97.1.298. [1]

Goolsbee, Austan, and Chad Syverson. 2021. “Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020.” *Journal of Public Economics* 193: 104311. DOI: 10.1016/j.jpubeco.2020.104311. [2]

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47 (2): 263–92. DOI: 10.2307/1914185. [1]

Kőszegi, Botond, and Matthew Rabin. 2006. “A Model of Reference-Dependent Preferences.” *Quarterly Journal of Economics* 121 (4): 1133–65. DOI: 10.1093/qje/121.4.1133. [1]

Krosnick, Jon A., and Duane F. Alwin. 1989. “Aging and Susceptibility to Attitude Change.” *Journal of Personality and Social Psychology* 57 (3): 416–25. DOI: 10.1037/0022-3514.57.3.416. [1]

Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics* 112 (2): 443–77. DOI: 10.1162/003355397555253. [1]

Sari, Berna A., Ernst H. W. Koster, and Nazanin Derakshan. 2017. “The Effects of Active Worrying on Working Memory Capacity.” *Cognition and Emotion* 31 (5): 995–1003. DOI: 10.1080/02699931.2016.1170668. [2]

Tversky, Amos, and Daniel Kahneman. 1974. “Judgment under Uncertainty: Heuristics and Biases.” *Science* 185 (4157): 1124–31. DOI: 10.1126/science.185.4157.1124. [1]

Chapter 1

Experienced Political Competition and Voter Apathy*

1.1 Introduction

Voting is foundational to the concept of democracy. By each casting their vote, citizens bring about an election outcome that is representative of their aggregated preferences, thereby granting legitimacy to the elected government. Yet, voter apathy is a widespread phenomenon in established democracies, which often exhibit low levels of political participation. This is especially true for the US, where voter turnout has consistently been well below 70 percent of the voting-eligible population even in presidential elections. Since people with low socioeconomic status go to the polls in particularly low proportion, this also tends to bias electoral participation towards the privileged, reinforcing concerns about the representativeness of election outcomes and resulting government policy (e.g., Lijphart, 1997; Fleck, 1999; Mueller and Stratmann, 2003; Brunner, Ross, and Washington, 2013). Consequently, understanding the driving forces behind voter apathy plays an essential role in efforts to build up and preserve a sustainable democratic society.

Combining insights from economic models of voting and social psychology, this chapter investigates growing up in an environment of low political competition as one potential determinant of voter apathy. This hypothesis builds on the key role that economic models of voting assign to the competitiveness of elections: the closer an election is expected to be, the higher is each individual voter's chance of affecting the outcome, implying an increase in the expected benefit of voting relative to its cost.[1] Similarly, political stakeholders like politicians and interest groups have a stronger incentive to mobilize potential voters in more competitive elections. However, while economic voting models limit the influence of electoral competition to the decision to vote in the current election, I lay the focus on its potential to affect individuals' propensity to vote in subsequent elections. This long-term perspective on the consequences of political

* I thank Peter Andre, Lorenz Goette, Luca Henkel, Paul Schäfer, and Florian Zimmermann for valuable comments and discussions. The ANES is based on work supported by the National Science Foundation under grant numbers SES 1444721, 2014–2017, the University of Michigan, and Stanford University.

1. See Feddersen (2004) for a review of the economic theory of the decision to vote or abstain in elections.

competition is inspired by the impressionable years hypothesis from social psychology, which posits that political values and attitudes underlying behavior are shaped during early adulthood (especially the age span of 18 to 25 years) and do not change much afterwards (e.g., Krosnick and Alwin, 1989). If the relevance of the impressionable years hypothesis extends to the act of voting itself, this would suggest that experiencing a lack of political competition during early adulthood could have a lasting negative impact on the propensity to vote later in life.

Conducting a convincing empirical test of this hypothesis is challenging because political competition experiences are correlated with a wide range of potential confounding factors. For instance, birth cohorts that live through a specific period of political competition also share many other experiences like economic crises or political scandals that might affect subsequent voting attitudes, and past political conditions in a given region might continue into the present. My identification strategy overcomes these hurdles by means of a fixed effects approach, using repeated cross-sectional survey data that spans over 60 years of elections within the unique electoral system of the US. Presidential elections in the US have historically been decided in a few battleground states, which receive the majority of attention during the campaign. This implies that the degree of political competition varies both across cohorts and states, allowing me to filter out unobservable confounders related to both of these dimensions by including fixed effects for an individual's birth cohort and the state they grew up in. In addition, fixed effects for the state of current residence interacted with election year capture the current electoral context, while age fixed effects capture systematic differences in the propensity to vote across the life cycle. Holding all else equal, my identification strategy then compares the self-reported turnout decisions of individuals who grew up in more or less competitive conditions, as measured by whether different or the same party won their home state in the two presidential elections experienced during the impressionable years.

My main finding is that individuals who grew up during periods of low political competition in their state are on average 1.4 percentage points less likely to vote later in life, which is in the range of causal estimates for other determinants of voting in the literature. Based on a linear approximation, the estimated effect is still significantly different from zero for individuals at the age of 45, implying that it endures for more than two decades on average. Moreover, the effect is stronger for individuals with lower education or lower income. These findings imply that periods of reduced political competition can have a long-lasting negative impact on future political participation and contribute to voter apathy. The resulting reductions in average turnout can weaken the legitimacy of elected officials and constrain policymaking. Since the effect is stronger for people with low socioeconomic status, it also promotes distortions in political outcomes and resulting policy choices relative to the democratic ideal of equal representation.

A causal interpretation of my results requires that there are no other unobserved determinants of voter apathy that simultaneously vary across birth cohorts and states and are correlated with political competition, since this is the dimension not covered by the fixed effects. To support this assumption, I conduct a variety of robustness checks. To begin with, I consider the possibility that my findings are driven by unobserved short-run factors that differentially affect wider geographic regions or by long-run changes in the propensity to vote that vary by state. I document that the effect persists in specifications that rule out these concerns by including additional interacted fixed effects for individuals' birth cohort by the census region they grew up in and for the state individuals grew up in by their 30-year cohort group. Then, I turn to specific time-varying state characteristics that might be correlated with the degree of political competition. In particular, low political competition could be accompanied by adverse economic and social conditions, which in turn might affect subsequent political participation. To alleviate this concern, I conduct regressions that control for measures of income and population growth, inequality, an indicator for experiencing a recession, the home ownership rate as a proxy for wealth, and the experienced crime rate, but the coefficient on experienced political competition barely changes. While these analyses cannot fully compensate for the lack of quasi-experimental variation in experienced political competition, they encourage a causal interpretation of the documented effect.

In the second part of the chapter, I investigate potential underlying mechanisms of my main result. In particular, I consider the possibility that political experiences affect subsequent voting inclination through the transmission of parental voting behavior, by encouraging the formation of a voting habit, or through a persistent increase in either the strength of individuals' political identity, the perceived civic duty to vote, or beliefs about the benefit of voting. This analysis is subject to the caveat that some of the survey items I use are only available in a small number of survey years, leaving me with a limited number of observations. I find that individuals who grew up under different levels of electoral competition do not differ in the political interest of their parents or their own strength of party identification, suggesting that the effect is not driven by an intergenerational transmission of parental voting behavior or an increase in the salience or development of people's political identity. Rather, individuals who experienced more competitive elections during their impressionable years seem to have more favorable attitudes towards voting later in life. For instance, they are more likely to agree that voting can make a difference, though the effect is only marginally significant. Therefore, one plausible interpretation is that people form enduring beliefs about the benefit of voting during their impressionable years, and these beliefs determine their propensity to vote in subsequent elections.

My findings belong to a strand of literature in political economy that investigates the relevance of electoral competition for political outcomes. On the topic of political selection, existing studies demonstrate that competition enhances the quality of elected politicians (De Paola and Scoppa, 2011) and reduces the likelihood of political dynasties (Dal Bó, Dal Bó, and Snyder, 2009). Competition is also an important requirement for political accountability, increasing the alignment between representatives' decisions and the policy preferences of their constituents (e.g., Balles, Matter, and Stutzer, 2022) as well as diminishing politicians' outside earnings, political rents and instances of corruption (Becker, Peichl, and Rincke, 2009; Svaleryd and Vlachos, 2009; Ferraz and Finan, 2011). Moreover, Besley, Persson, and Sturm (2010) show how a lack of political competition can lead to the adoption of government policies that hinder economic growth. My findings contribute to this line of research by demonstrating that low political competition also harms the democratic system by entrenching voter apathy among the electorate, thus threatening the representativeness of future election outcomes and implemented public policy. The uncovered heterogeneity of this effect by socioeconomic status resonates with recent studies that point to additional hurdles in the voting process for disadvantaged groups, like longer wait times at polling places and a higher incidence of voting errors (Shue and Luttmer, 2009; Chen, Haggag, Pope, and Rohla, Forthcoming).

More generally, my study is related to an emerging literature on the effects of personal experiences on political attitudes and voting behavior.[2] Alesina and Fuchs-Schündeln (2007) analyze the effect of growing up in a communist regime on a range of policy preferences and beliefs. Several studies show that redistributive preferences and vote choice are affected by recession and inequality experiences during the impressionable years (Giuliano and Spilimbergo, 2014; Roth and Wohlfart, 2018; Carreri and Teso, Forthcoming). There is also evidence regarding the role of experiences for the level of political participation. For instance, support for democracy is partly determined by exposure to a democratic system (Fuchs-Schündeln and Schündeln, 2015; Acemoglu, Ajzenman, Aksoy, Fiszbein, and Molina, 2021). Moreover, Akbulut-Yuksel, Okoye, and Yuksel (2020) and Alacevich and Zejcirovic (2020) show in two historical contexts that violence against civilians during war and witnessing expulsions of other ethnic groups persistently decreased voter turnout in subsequent elections. Similar effects are observed by Eichengreen, Saka, and Aksoy (2022) for individuals who experience an epidemic. I extend this literature by documenting that not only drastic experiences like war, autocracy and epidemics can be a threat to political participation. Even variation in political competition within a functional democratic system like the US can have a persistent effect on people's propensity to abstain from voting.

Most empirical research on the determinants of voter turnout is concerned with concurrent influences on the propensity to vote, like the institutional context of the election or mobilization activities by political parties.[3] With respect to electoral competition, my findings indicate that this focus on the present might be too narrow. Considering the possibility of a persistent effect of past experiences allows to put a different interpretation on the results of recent studies that test whether voter turnout is increasing in the perceived closeness of the election, as implied by many economic models of voting. While abstract laboratory experiments mostly support this prediction (e.g., Levine and Palfrey, 2007; Faravelli, Kalayci, and Pimienta, 2020), Gerber, Hoff-

2. Experience effects are also pervasive in many other domains of economic behavior: they have been shown to affect market beliefs (Di Tella, Galiani, and Schargrodsky, 2007), financial risk taking and stock market participation (Malmendier and Nagel, 2011), inflation expectations (Malmendier and Nagel, 2016; Malmendier, Nagel, and Yan, 2021), consumption patterns (Malmendier and Shen, 2018), preferences for driving (Severen and van Benthem, 2022), and the emancipation of women (Slotwinski and Stutzer, 2018).

3. For instance, some recent studies examine the effects of media market expansions and technological advances which change the information consumption of eligible voters (e.g., Gentzkow, 2006; Enikolopov, Petrova, and Zhuravskaya, 2011; Gentzkow, Shapiro, and Sinkinson, 2011; Falck, Gold, and Heblich, 2014), labor market conditions (Charles and Stephens Jr., 2013), institutional factors that affect the cost of voting (Hodler, Luechinger, and Stutzer, 2015; Cantoni, 2020; Kaplan and Yuan, 2020; Cantoni and Pons, 2021), canvassing and political advertising (Pons, 2018; Spenkuch and Toniatti, 2018; Pons and Liegey, 2019), and the role of social image concerns (DellaVigna, List, Malmendier, and Rao, 2017) and expressive voting (Pons and Tricaud, 2018). More closely related to my work are studies that share my focus on the relevance of past circumstances for current political participation. Among these, Fujiwara, Meng, and Vogl (2016) and Bechtel, Hangartner, and Schmid (2018) explore the role of habit formation, which is one potential mechanism underlying my results. Other contributions trace contemporary patterns in political participation to historical acts of racial violence (Williams, Forthcoming) or investigate the long-term effects of changes in the sociodemographic environment during childhood (e.g., Chyn and Haggag, 2019; Cohodes and Feigenbaum, 2021).

man, Morgan, and Raymond (2020) find little evidence of a causal effect of pivotality beliefs on the intention or decision to vote in two field experiments, casting doubt on the importance of perceived competitiveness in the context of large real-world elections.[4] My results on mechanisms suggest that to some extent, attitudes towards the benefit of voting depend on the degree of political competition experienced during the impressionable years rather than the expected closeness of the current election. These persistent attitudes might then drive the propensity to vote in actual elections, whereas they might not be activated in abstract laboratory games.

The idea that the political environment during early adulthood shapes subsequent voting behavior has also been discussed in political science. Building on evidence for the persistence of turnout across the life cycle (e.g., Plutzer, 2002), Franklin (2004) evokes changes in the electoral competition experienced by newly eligible cohorts in combination with habit formation as a potential explanation for long-term aggregate turnout dynamics. However, his own empirical analysis only includes political competition as a short-term factor influencing current turnout levels.[5] Subsequent empirical studies have found mixed results, with three articles arguing for (Pacheco, 2008; Smets and Neundorf, 2014; Vowles, Katz, and Stevens, 2017) and one against (Heath, 2007) a persistent effect of the experienced electoral context on future turnout.[6] These studies typically either do not consider the possibility of unobservable confounding factors due to common region-, cohort-, and time-specific shocks or rely on hierarchical models with random coefficients to account for such group-specific effects. While this has the advantage that the group-specific effects can be directly estimated, the coefficient of interest will be biased if there is a correlation between a predictor and the group effects (e.g., Bafumi and Gelman, 2006).[7] In contrast, my fixed effects approach in combination with cross-sectional and cohort-level variation in experienced political competition alleviates such endogeneity concerns, thus encouraging a causal interpretation of the results.

The remainder of the chapter is structured as follows. In Section 1.2, I briefly outline my hypothesis about the effect of experienced political competition on individuals' subsequent propensity to vote in the context of US presidential elections. I describe the data and the em-

4. Observational evidence in favor of an effect of anticipated closeness on voter turnout comes from studies that rely on plausibly exogenous variation in exposure to pre-election or exit polls (Morton, Muller, Page, and Torgler, 2015; Bursztyn, Cantoni, Funk, Schönenberger, and Yuchtman, 2020). However, it is difficult to fully rule out that these effects are driven by reactions of political elites or changes in the perceived social pressure to vote rather than pivotality beliefs alone.

5. In general, the most commonly used strategy to assess Franklin's theory has been to test whether measures of current political competition have a stronger effect on turnout for newly eligible as compared to older cohorts, without an investigation of whether these effects are enduring (e.g., Franklin, Lyons, and Marsh, 2004; Smets, 2012; Blais and Rubenson, 2013; Górecki, 2013).

6. Most closely related to my work, Smets and Neundorf (2014) analyze the experience effects of several national-level measures of the electoral context in the US, using data from the General Social Survey. While the closeness of the first few experienced elections does not have an effect on turnout in their hierarchical model, they find that experiencing a higher aggregate turnout rate has a small positive impact on the decision to vote later in life. In addition to the differences in the econometric approach, I also go beyond the scope of their study by considering experienced political competition at the state-level, which I find to be the more important factor in the context of the US.

7. See Bell and Jones (2015) for a thorough discussion of the advantages of random effects models from a political science perspective.

pirical strategy used to test the developed hypothesis in Section 1.3 and Section 1.4, followed by the presentation of the main results in Section 1.5. In Section 1.6, I provide some suggestive evidence on the underlying mechanisms of the uncovered effect. Finally, Section 1.7 concludes.

1.2 Hypothesis in the Context of the US

Social psychologists take the view that people's basic social and political attitudes are formed during adolescence and early adulthood and stabilize with age. One particularly influential theory is the impressionable years hypothesis, which puts a special emphasis on the age between 18 and 25 for the development of various aspects of the value system (e.g., Krosnick and Alwin, 1989; Sears and Funk, 1999). Experiences during this period of life are said to leave an imprint that can shape behavior for many years to come. This may be reducible to neurochemical and anatomical changes in the adolescent brain that yield a heightened mnemonic capacity during this period, implying that memories of experiences made during late adolescence and early adulthood are more likely to be recalled than memories from other stages of life (Spear, 2000; Fuhrmann, Knoll, and Blakemore, 2015).

In most democratic countries, it is exactly this age span in which people start to participate in the political process and can thus be expected to form initial political opinions. Around that time, they become eligible to vote, are targeted by party mobilization efforts and political advertising, and experience their first elections from the perspective of a voter. Hence, it is not surprising that the impressionable years hypothesis has generally found empirical support in the domain of political preferences. For instance, Ghitza, Gelman, and Auerbach (2021) show that the presidential approval ratings experienced during early adulthood predict more than 80 percent of the variation in voting trends in the US over the last half-century. Moreover, recent studies in economics have found that recessions and the level of inequality experienced during the impressionable years determine preferences for redistribution and vote choice (Giuliano and Spilimbergo, 2014; Roth and Wohlfart, 2018).

In light of this evidence, it seems plausible that the relevance of the impressionable years hypothesis extends to the act of voting itself. This would imply that individuals' election experiences during early adulthood determine whether they develop into politically participating citizens or slip into voter apathy. According to many economic models of voting, one key aspect of these election experiences is the prevailing degree of political competition, which controls both each individual voter's chance of affecting the election outcome as well as the incentives of candidates and interest groups to invest in voter mobilization (e.g., Downs, 1957; Palfrey and Rosenthal, 1983; Ledyard, 1984; Shachar and Nalebuff, 1999). In some models, ethical voters also follow endogenous ethical rules or voting norms that depend on the competitiveness of the election in equilibrium (e.g., Feddersen and Sandroni, 2006; Ali and Lin, 2013).

The Electoral College system of the US vividly illustrates the role of electoral competition for individuals' election experiences. In presidential elections, each state allocates a fixed number of Electoral College votes to one of the presidential candidates, and the candidate who receives the majority of these votes wins the election.[8] The electoral votes of each state are

8. The number of electoral votes of each state reflects its population size.

awarded according to a winner-takes-all system: the presidential candidate who receives the majority of the popular vote in a given state receives all its electoral votes, irrespective of the margin of victory. As a consequence, the attention of the media and the presidential campaigns is focused on a few highly competitive "battleground states" where the election is expected to be close, while the residents of less competitive states only take on a spectator role. In particular, Shaw (1999) and Strömberg (2008) show that presidential campaigns mostly target battleground states in their use of resources and planning of campaign activities.[9] This includes rallies and speeches by presidential candidates and other high-profile politicians, political advertisements, and door-to-door canvassing aimed at mobilizing supporters to vote. All of these factors presumably have a positive effect on political efficacy beliefs, the development of a political identity, and the perceived civic duty to vote that is only reaped by residents of more competitive states.

In addition to these differences in external influence factors, experiences also diverge on a psychological level. In spectator states, the winner of the popular vote is certain long before Election Day, making residents feel as if their votes were irrelevant for the election outcome. This is also reflected in the political atmosphere in the build-up to the election, with lopsided pre-election polls and few incentives to discuss politics with neighbors or engage in campaign activities oneself (Settle, Bond, Coviello, Fariss, Fowler, et al., 2016).

In sum, growing up in less competitive states leads to election experiences that—due to a variety of factors—seem less conducive to political participation. The impressionable years theory posits that these experiences, when made during the age span of 18 to 25, leave a lasting imprint on the young brain that could differentially affect voting behavior even when the level of political competition has converged again. Thus, my main hypothesis is that individuals who grow up in less competitive states will have a lower propensity to vote later in life, holding constant other determinants of voting.

1.3 Data

Investigating the effect of personal experiences of state-level electoral competition during the impressionable years requires individual-level data on turnout, birth cohort and the state one grew up in. While administrative turnout data would be preferable, data on voters' state of residence during early adulthood is typically not included in state voter files, so I have to rely on survey data.[10] However, one advantage of using survey data is that it also covers eligible voters who are not registered to vote (Jackman and Spahn, 2021), especially since the decision to register might also be affected by political experiences.

My main data source is the American National Election Studies (ANES). The ANES is a repeated cross-sectional survey of the voting-eligible population in the US that has been conducted in most midterm and presidential election years since 1948. In each survey year, a

9. Strömberg (2008) also shows that this behavior is a best response to the incentives of the Electoral College in a probabilistic-voting model of political competition.

10. The downside of self-reported turnout data is that it is often subject to substantial overreporting, presumably because nonvoters do not like to admit that they did not vote (DellaVigna et al., 2017).

nationally representative sample of about 1000 to 2500 respondents is interviewed about their political views and election behavior. The basis for my analysis is the ANES Time Series Cumulative Data File (The American National Election Studies, 2019b), which I supplement with information about the sampling design of each survey from the individual year ANES Time Series Studies (The American National Election Studies, 1999a,b,c,d,e,f,g,h,i,j,k,l,m,n,o,p,q,r,s, 2002, 2003, 2005a,b, 2015, 2016, 2019a). I also use data on the state respondents grew up in from the 2004, 2012 and 2016 individual surveys because this data is not yet included in the cumulative data file.

In the analysis, I pool presidential and midterm elections, so my primary outcome of interest is an indicator variable equal to one if the respondent reports having voted in the respective year's national election (item VCF0702). In additional analyses, I also use survey responses to questions about political identity and attitudes towards voting as dependent variables.

1.3.1 Construction of Political Competition Experiences

To calculate eligible voters' political competition experiences, I collect state-level election results (including Washington, DC) for all presidential elections between 1900 and 2012 from the American Presidency Project (Woolley and Peters, 2019). I focus on presidential elections because this is the most important type of election in the US, so I expect experiences from these elections to have an especially strong impact. Since presidential elections are held every four years, each birth cohort experiences exactly two presidential elections during the impressionable years period between age 18 and 25.[11]

I measure the experienced political competition of birth cohort c growing up in state s by constructing an indicator variable equal to one if the same party won their state in both presidential elections held during their impressionable years, such that a value of one indicates low experienced competition.[12] The intuition behind this choice is that a change in the winning party implies that either at least one of the two elections was contested or there was a reversal in the political atmosphere within the state. In both cases, it seems likely that there was a high level of campaign intensity in at least one election. On a psychological level, experiencing a change in the winning party should additionally convey the impression that voting can make a difference to the election outcome.[13]

Figure 1.1 shows the variation in experienced political competition across states and election cohorts, defined as the year of the first presidential election a birth cohort experiences in

11. There was a change in the legal voting age from 21 to 18 years that first came into effect for the majority of states in the presidential election of 1972. This implies that while everyone experiences two presidential elections during their impressionable years, some respondents in my sample were only eligible to vote in one of them. The potential effect of this legislative change on turnout is accounted for by my identification strategy.

12. Since the US has a two-party system dominated by the Republican and Democratic parties, I ignore third parties in the construction of the measure.

13. An alternative measure of political competition could be based on the average state-level victory margin of the two elections. However, this has the disadvantage that the average margin of victory can be high even if one of the elections is fiercely contested. Moreover, it is unclear what functional form the relationship between turnout and this measure should be expected to take, as it is unlikely to be linear. Nonetheless, I also explore this possibility in Section 1.5.4.

their impressionable years.[14] Evidently, there is a lot of variation both within states across time and across states for a given cohort.

I merge political experiences to respondents of the ANES sample on their birth cohort and the state they grew up in. The geographic question (items VCF0132 and VCF0133) asks about the state respondents' grew up in without specifying an age range.[15] Presumably, interviewees think about their childhood when answering rather than naming their place of residence during their impressionable years. Thus, I only measure the actual political experiences with error because some respondents might move to another state to study or work after finishing school. However, since measurement error in an independent variable biases the regression coefficient towards zero, this reduces the probability of falsely rejecting the null hypothesis of no effect. Moreover, there is evidence that the majority of Americans never move away from their home state (Taylor, Morin, Cohn, and Wang, 2008).

1.3.2 Sample

Not all respondents for whom political experiences can be constructed end up in the estimation sample. First, I exclude individuals who are age 25 or younger at the time of the ANES interview because my hypothesis concerns turnout decisions after people have lived through their impressionable years. I also exclude individuals who are 90 years old or above because the available age variable is truncated at this value, preventing me from matching political competition experiences to these respondents. Moreover, I restrict the sample to respondents who have had voting rights in their state of residence during their impressionable years. This precludes female respondents who grew up before women's suffrage was introduced in their state, and cohorts from Washington, DC, who experienced a presidential election during their impressionable years before the federal district was granted voting rights in 1961.[16] I also drop cohorts from Nebraska and Maine who experienced an election during their impressionable years after these states changed the allocation rules for their electoral votes from a winner-takes-all to a split system, which makes it more difficult to assess the level of political competition.[17] Lastly, the ANES cumulative file includes a small number of panel study respondents who were interviewed in more than one election year. I exclude all but the first occurrence of these respondents in the dataset.

The final sample consists of roughly 32 800 respondents and covers 26 US elections between 1952 and 2016. Descriptive statistics of the most important variables for the unweighted sample

14. Since presidential elections are held every four years, each election cohort comprises four birth cohorts. For example, election cohort 2000 consists of birth cohorts 1979 to 1982. In contrast, someone born in 1978 experienced the first presidential election of their impressionable years at age 18 in 1996, and someone born in 1983 experienced the first presidential election of their impressionable years at age 21 in 2004.

15. The question allows for multiple mentions. I drop respondents who name more than one state.

16. Data on the state-level introduction of women's suffrage was taken from Table 1 in Lott Jr. and Kenny (1999).

17. These states now allocate two electoral votes to the winner of the state-level popular vote and one electoral vote to the winner of each congressional district within the state. Maine implemented this rule prior to the presidential election of 1972 and Nebraska before the election of 1992, so the affected birth cohorts are 1947 and later in Maine and 1967 and later in Nebraska.

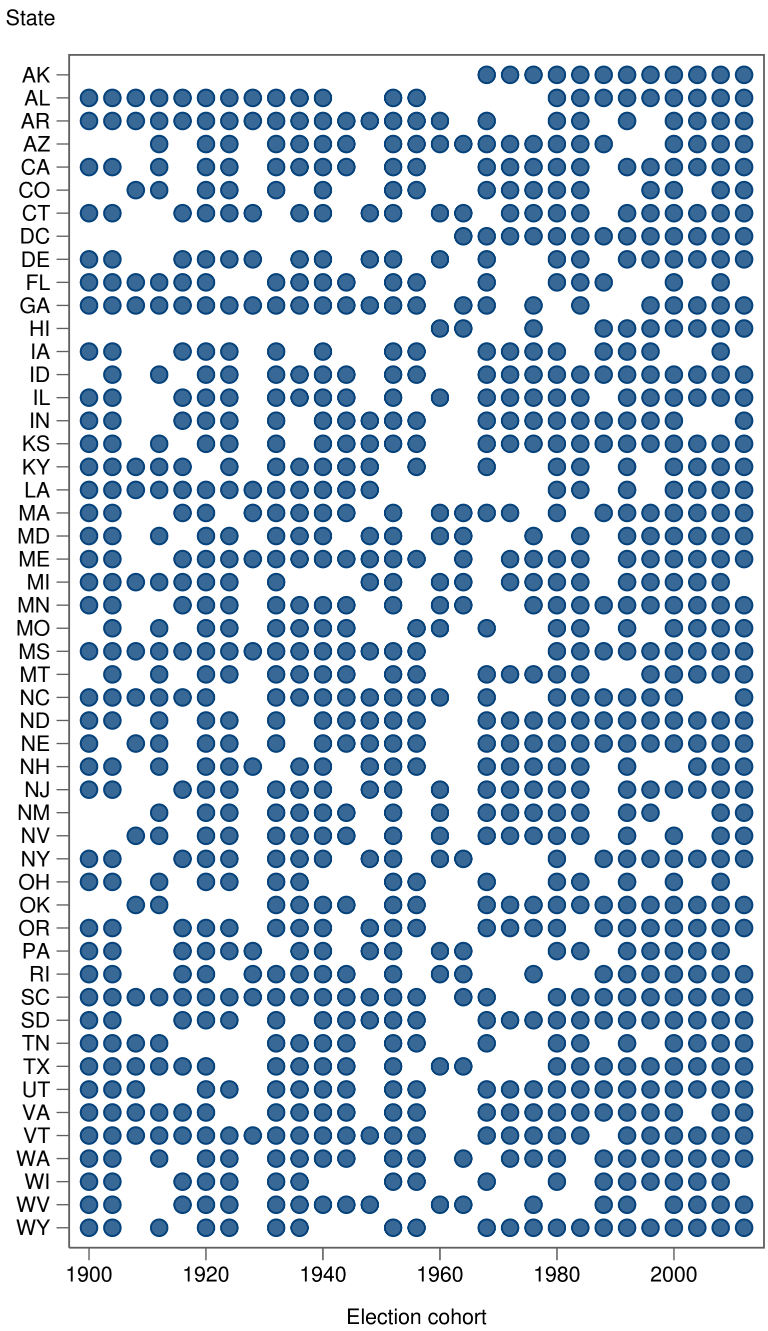

Figure 1.1. Experienced Political Competition across Election Cohorts and States

Notes: Dots indicate low experienced political competition. Each election cohort comprises the four birth cohorts who experienced the first presidential election of their impressionable years in the given year. A given election cohort who grew up in a given state is classified as having experienced low political competition if that state was won by the same major party in both presidential elections during that cohort's impressionable years.

are reported in Table 1.B.1. Subsequently, all analyses use the provided ANES survey weights to make the respondents from each survey year representative of the US voting-eligible population at that time.

1.4 Empirical Strategy

I want to compare the current turnout decisions of individuals who differ in the political competition they experienced during their impressionable years, but are alike in all other respects. To achieve this, I rely on an empirical strategy that is similar to previous economic studies of experience effects (Giuliano and Spilimbergo, 2014; Roth and Wohlfart, 2018). The basic intuition is that my setting permits a fixed effects approach that allows to flexibly control for the major dimensions of heterogeneity in the data.

The key advantage of my setting is that political experiences vary both across election cohorts within the state respondents grew up in and across these states within election cohorts, so the effect of experienced political competition on turnout is identified even when fixed effects for both of these dimensions are included in the regression.[18] At the same time, political experiences also vary within group cells for all other major dimensions of heterogeneity in the data. This is obvious for some dimensions, e.g., that political experiences vary between different respondents surveyed in a given election year, and a result of the context of my study in others, like the variation of my measure of political competition within respondents of the same age. As a result, I can flexibly control for many unobservable confounding factors by including several dimensions of fixed effects in the estimation, as long as the group-specific effects are additively separable. The effect of experienced political competition is still identified by the remaining within-group variation.

The estimation equation is

$$\begin{aligned} Voted_{ist} \;=\; & \beta SameStateWinner_i + \boldsymbol{\gamma}'\boldsymbol{X}_{it} + \eta_s \times \delta_t + \alpha_{age} + \kappa_{cohort} + \rho_{state\text{-}when\text{-}young} \\ & + \epsilon_{ist}, \end{aligned} \tag{1.1}$$

where $Voted_{ist}$ is an indicator equal to one if individual i, living in state s, voted in election year t. $SameStateWinner_i$ is an indicator variable measuring i's experienced political competition during the impressionable years as outlined in Section 1.3.1, such that a value of one corresponds to low experienced competition. Thus, β is the coefficient of interest, and my hypothesis corresponds to $\beta < 0$. To accommodate the use of multiple levels of fixed effects and for ease of interpretation, I estimate Equation 1.1 as an ordinary least squares (OLS) regression, using the estimator developed in Correia (2016).

The vector $\boldsymbol{X}_{it}$ includes control variables for a range of demographic characteristics that previous research has shown to predict turnout (e.g., Ashenfelter and Kelley Jr., 1975; Glaeser, Ponzetto, and Shapiro, 2005; Leighley and Nagler, 2013). In particular, I add dummies for education level, marital status, religiosity, and four different relative income groups. In addition,

18. Remember that each election cohort comprises the four birth cohorts that experience the same two presidential elections during their impressionable years.

I also control for gender, ethnicity, and employment status.[19] I do not include a measure for the strength of party identification even though it predicts turnout because party identification might itself be affected by experienced electoral competition.

The remaining parameters $\eta_s \times \delta_t$, α_{age}, κ_{cohort}, and $\rho_{state\text{-}when\text{-}young}$ denote fixed effects for the current state of residence interacted with election year, age, birth cohort, and the state a respondent grew up in, respectively.[20] State-year fixed effects capture all effects of the current electoral context on turnout, like the current degree of political competition, differences in voting rules, or the appeal of a running candidate, which may vary by state.[21] Among other things, this rules out that the estimate of β is biased due to serial correlation in political competition on the state-level. Life cycle effects in the propensity to vote are accounted for by age fixed effects. This is important because age could be correlated with political experiences through, for instance, a time trend in electoral competition. Birth cohort fixed effects capture other formative experiences common to a cohort that might affect turnout and be correlated with experienced political competition. For instance, experiencing the Vietnam War might lead to distrust towards the government and increase political polarization. Finally, state-when-young fixed effects account for time-invariant differences in the propensity to vote across respondents' home states, like historical variation in instilled voting norms.

I cannot include fixed effects for cohort interacted with state-when-young because this is the level of variation of the regressor of interest. Thus, an important identification assumption is the absence of unobserved heterogeneity in turnout that simultaneously varies across birth cohorts and states and is correlated with political competition, conditional on the control variables. Intuitively, differences in later-life turnout between cohorts across states must only be due to changes in political competition or due to factors that are orthogonal to political competition. While this assumption is untestable, I conduct several analyses to gauge its validity in Section 1.5.3.

In a second specification, I also investigate whether there is an effect of national-level political competition experiences on future turnout, intended as a rough comparison to the effect of state-level experiences. For this purpose, I replace $SameStateWinner_i$ with a similarly constructed indicator that is equal to one if the same party wins the presidency in both presidential elections experienced during a respondent's impressionable years. Since there is no cross-sectional variation in these experiences, this specification uses cohort group fixed effects

19. To avoid a loss of statistical power due to missing values on control variables, each category also includes a dummy equal to one if data for that category is missing for the respondent, as in Roth and Wohlfart (2018). The coding of all control variables is described in detail in Appendix 1.A.

20. Year, birth cohort, and age effects are not separately identified because the linear parts of these time effects are collinear, i.e., it cannot be disentangled whether an effect is due to age, cohort, or year. This is not a problem for my identification strategy because I am not interested in estimating or interpreting these effects, but only use them to filter out unobserved confounding factors that would otherwise bias the coefficient of interest.

21. For instance, Cantoni and Pons (2022) estimate that on average, holding fixed a given election, contextual factors of eligible voters' current state of residence explain about one-third of the observed variation in turnout between groups of states.

instead of birth cohort fixed effects.[22] Thus, it is vulnerable to confounding factors that are specific to certain birth cohorts, but not to long-term changes in voting attitudes across generations.

In all regressions, I use two-way cluster-robust standard errors to correct for both correlated treatment assignment and complex survey sampling at the same time (Cameron, Gelbach, and Miller, 2011; Abadie, Athey, Imbens, and Wooldridge, 2017). Since political competition experiences are the same for people who are born in the same state within the same election cohort, the first level of clustering is by election cohort interacted with state-when-young. The second dimension of clustering is by primary sampling unit (PSU) of the respective ANES survey. This accounts for the fact that only a geographically clustered subset of the whole population is interviewed in each survey year. When constructing the cluster groups for this dimension, I pool respondents from the same PSU of different ANES waves that rely on the same sampling frame, while treating survey years that use different survey designs as independent samples (Korn and Graubard, 1999).[23]

1.5 Results

In this section, I present the results of my analysis. I start off with some preliminary graphical evidence and then go on to report the results of the main test of my hypothesis. Afterwards, I discuss the plausibility of a causal interpretation of the result and briefly summarize the robustness checks I conducted. Next, I investigate heterogeneity by socioeconomic status and time passed since the impressionable years. Finally, I report on some additional findings.

1.5.1 Preliminary Evidence

To convey an impression of the data, I first present some preliminary graphical evidence. Figure 1.2 displays the average residuals from a regression of an indicator for voting on state interacted with election year fixed effects, plotted separately for individuals who are classified as having experienced low and high electoral competition. The bars show that a lack of political competition during early adulthood is associated with a lower propensity to vote at the time of the ANES survey. The raw difference in the likelihood of voting between the two groups that cannot be explained by the current electoral context amounts to about two percentage points. However, this raw difference could also be spuriously driven by a correlation of political experiences with other determinants of voting like age or cohort effects, which motivates the remainder of my analysis.

22. Cohort groups are defined as 30-year intervals over the approximate range of birth cohorts from the main estimation sample: 1871 to 1900 (defined as 1900 and earlier), 1901 to 1930, 1931 to 1960, and 1961 to 1990 (defined as 1961 and later).

23. Details regarding the construction of the cluster groups are described in Appendix 1.A.

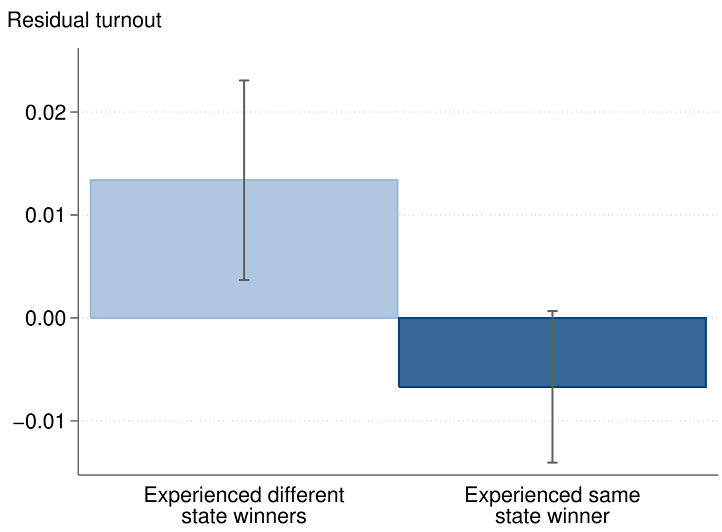

Figure 1.2. Residual Turnout by Experienced Political Competition

Notes: Bar chart of the average residuals from a regression of an indicator for voting in the current election on state interacted with election year fixed effects for respondents who experienced high political competition (left bar) and low political competition (right bar).

1.5.2 Main Results

Estimates of the main specification described in Section 1.4 are presented in column (1) of Table 1.1. The coefficient of interest is negative and significantly different from zero, lending support to the tested hypothesis. Respondents who experienced low political competition in their home state during the impressionable years are 1.4 percentage points less likely to vote in subsequent elections than those who experienced a more competitive political environment, even after accounting for differences in the current electoral context as well as arbitrary heterogeneity by birth cohort, age, and state-when-young. This indicates that a lack of political competition can have a negative impact on future political participation and contribute to voter apathy.

To get a sense of the magnitude of the effect, I compare it to causal estimates of other determinants of voting from the literature. Chyn and Haggag (2019) estimate that moving children from disadvantaged areas to more affluent neighborhoods increases their probability of voting in a given presidential election as adults by about 2.3 to 2.8 percentage points. Since such a move is associated with improvements in various later-life outcomes that have been linked to political behavior, including education, incarceration rates, and labor market success, this comparison suggests that the effect of experienced electoral competition is sizable. In terms of voting costs, it corresponds to a reduction in travel distance to the polls of 0.1 to 0.3 miles (Cantoni, 2020) or to the provision of seven additional days of early voting—an increase of 50 percent relative to the current average number of early voting days across US states (Kaplan and Yuan, 2020).

Table 1.1. Main Results

	Dependent variable: Voted in last election		
	(1)	(2)	(3)
Experienced same state winner	−0.014**		−0.011*
	(0.006)		(0.006)
Experienced same national winner		−0.005	−0.002
		(0.005)	(0.006)
Demographic controls	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Birth cohort FE	Yes	No	No
Cohort group FE	No	Yes	Yes
State-when-young FE	Yes	Yes	Yes
Observations	32 768	32 768	32 768
R^2 (within)	0.073	0.074	0.075
Mean (dependent variable)	0.714	0.714	0.714

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced same national winner* is an indicator equal to one if the presidency was won by the same major party in both presidential elections experienced during the impressionable years. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

In column (2), I replace my main measure of low experienced political competition with its national-level equivalent. With the caveat that other differences across birth cohorts might also play a role in this specification, the effect of experiencing low political competition on the national level is much smaller and not significantly different from zero. Including both state-level and national election experiences in the same regression, as in column (3) of Table 1.1, corroborates this conclusion. While the coefficient on state-level experiences is negative and remains significant at the ten percent level, the estimate of the effect of national experiences moves even closer to zero. Thus, it seems that the degree of political competition individuals experience in their own state is more important than the political environment in the nation as a whole. This result is not without intuitive appeal. Psychologically, the value of one's own vote in a spectator state might feel low precisely because of the concurrent existence of battleground states that are proclaimed by the media as the places where the election will be decided. In contrast, the differential competitiveness of previous presidential elections on the national level is less salient for voters and can easily be attributed to fluctuations in candidate appeal rather than a systemic irrelevance of the own vote. Moreover, even weak candidates spend resources on campaigning and voter mobilization in the states where they are most likely to win. Since the state rather than the nation seems to be the relevant context for experienced political competition, I focus on state-level experiences in the remainder of the analysis.

1.5.3 Causal Interpretation of the Main Results

A causal interpretation of the results requires that there are no unobserved determinants of voter turnout that simultaneously vary across birth cohorts and states and are correlated with political competition. This assumption is untestable, but I conduct several additional analyses to provide supportive evidence.

As a first step, I estimate the model in three additional fixed effects specifications that put an even tighter control on unobservable confounders. The results are presented in the first three columns of Table 1.C.1 in Appendix 1.C. In column (1), I interact state-when-young fixed effects with cohort group fixed effects to control for long-run changes in the propensity to vote that vary by state. The coefficient on experienced political competition barely changes. In column (2), I instead control for unobserved short-run factors that differentially affect wider geographic regions by interacting birth cohort fixed effects with fixed effects for the census region the respondent grew up in. In this specification, the magnitude of the effect even increases to a 2.2 percentage points lower likelihood of voting for individuals who experienced a lack of electoral competition. In the last column, I add both fixed effects interactions together in one specification. Again, the estimated coefficient on political experiences is negative and highly significant. This alleviates concerns that the results are driven by state-varying generational changes in zeitgeist or region-specific cohort effects.

In a second step, I control for a large set of time-varying historical state characteristics that might be correlated with experienced political competition and could potentially affect the later-life turnout of residents.[24] The coefficient estimates from these regressions are depicted in Figure 1.C.1.

One potential concern is that competition-reducing factors also adversely affect the economy or lead to higher inequality on the state-level, and this in turn reduces satisfaction with the government and political participation.[25] For instance, Besley, Persson, and Sturm (2010) show that low political competition can facilitate government policies that hinder economic growth. To rule out that my findings capture a response to the negative side effects of low political competition rather than to election experiences themselves, I control for various measures of social and economic conditions. In particular, I include state-level measures of income and population growth, inequality, and an indicator for experiencing a recession during one's impressionable years. Furthermore, I add the experienced home ownership rate in the state respondents grew up in to proxy for wealth.

I also run regressions that contain measures of property crime and violent crime. Blanco and Ruiz (2013) show that feeling insecure due to high levels of crime is related to lower levels of satisfaction with democracy. If their finding arises from a persistent causal effect, one might expect experienced crime to reduce subsequent voter turnout. At the same time, a high crime rate could affect concurrent political competition if one party is associated with a higher competence in imposing law and order. Moreover, high crime rates can be interpreted as a proxy for low social capital (Buonanno, Montolio, and Vanin, 2009), which has also been

24. The construction of all control variables is described in detail in Appendix 1.A.

25. However, Charles and Stephens Jr. (2013) find a negative relationship between local labor market conditions and voter turnout.

associated with reduced political participation. For instance, Akbulut-Yuksel, Okoye, and Yuksel (2020) show that witnessing Jewish expulsions in Nazi Germany decreased voter turnout in subsequent elections and explain this with a decrease in social capital.

My results are robust to the inclusion of all these historical control variables. While the coefficient on experienced political competition is not always statistically significant because data limitations in the control variables lead to a large decrease in sample size in some specifications, the point estimate consistently ranges between 1.2 and 1.5 percentage points.

1.5.4 Robustness

The results of additional robustness checks of the main results are reported in the remainder of Appendix 1.C. First, I consider the effect of experiences that are constructed from alternative measures of political competition (Table 1.C.2). One concern about the measure of electoral competition used in the main specification is that it overstates the incidence of low competition for respondents who experience two close elections with the same winner. As a robustness check, I instead use the absolute value of the victory margin in the state the respondent grew up in, averaged over the two presidential elections experienced during the impressionable years. Since the effect of political competition might be non-monotonic, I employ a flexible functional form by including indicator variables for each quartile of the distribution of the measure in my estimation sample. The estimated coefficients reveal that, compared to respondents whose experienced average victory margin lies in the lowest quartile, individuals who experienced higher victory margins are 1.2 to 1.6 percentage points less likely to vote in subsequent elections.[26] This effect size is in line with that estimated using my main measure of experienced political competition, though it is only significant for one of the three quartiles.[27] In the same table, I also show that accounting for third parties, which have won single states in presidential elections in rare cases, does not affect the results obtained in the main specification.

Second, I rerun the main specification with fixed effects for congressional district interacted with survey year, i.e., at a lower level of spatial aggregation (column (4) of Table 1.C.1). This addresses the concern that my sample also covers midterm elections, in which not all states hold state-wide races. About one-third of states in each midterm election do not have a senator up for election, so the only national race for parts of the electorate is for the House of Representatives. Moreover, even in presidential elections or if a Senate race is held, some eligible voters may care more for their district's representative in the House. Thus, the relevant current electoral context for some eligible voters might be the respective congressional district rather than the state. It is reassuring that the results remain unchanged when this possibility is accounted for.

Third, I verify that my findings retain their statistical significance under different assumptions about the clustering of standard errors. All main text regressions are estimated with two-way cluster-robust standard errors to account for both clustered treatment status and complex

26. The cutoff value between the lowest and second-lowest quartile is at an average victory margin of about six percentage points.

27. Using the same approach for national-level political experiences yields mixed results. There is some evidence of a negative effect of medium values of the victory margin for electoral votes, but not for the popular vote.

survey sampling. The estimates in Table 1.C.3 show that the standard errors barely change if I cluster in each of the two dimensions individually.

In Table 1.C.4, I report the estimated coefficients of specifications that either only control for the fixed demographic characteristics gender and minority status or additionally control for party identification by including indicator variables for identifying as a Democrat or Republican. The first of these specifications alleviates concerns that other demographic characteristics might themselves be affected by the degree of political competition experienced while growing up. The second specification shows that my results also hold within party. I also run regressions that exclude respondents who either grew up in the political south of the US or lived there in the year the survey was conducted.[28] These southern states went through a period of especially low electoral competition and voter turnout starting in the 1890s because Democratic state governments introduced various voting restrictions that effectively disenfranchised Black and poor voters (e.g., Besley, Persson, and Sturm, 2010; Naidu, 2012). The voting restrictions were abruptly abolished after the passage of the Voting Rights Act in 1965, resulting in a restoration of political competition and higher turnout rates. Thus, one might be concerned that other social or cultural changes in the southern states might partly drive the results. However, the estimated coefficient on low experienced political competition is always similar to that in the main specification.

1.5.5 Heterogeneity

From a normative perspective, reductions in the propensity to vote are especially problematic if they lead to unrepresentative electoral participation (e.g., Lijphart, 1997). If nonvoters were representative of the population at large, low voter turnout would cast doubt on the validity of popular consent to the elected government, but at least not have immediate consequences for election outcomes. If, however, the policy preferences of voters and nonvoters differ, low turnout also threatens democratic legitimacy by inducing violations of the principle of equal representation.

In reality, voter turnout in the US is vastly unequal. For instance, ANES respondents with a college degree were approximately 26 percentage points more likely to report having voted in the 2016 national election than respondents who did not complete high school.[29] Similarly, the gap in self-reported turnout between respondents in the top and bottom tercile of household income was about 22 percentage points in the same election.

Unequal voter turnout has consequences for government policy. Brunner, Ross, and Washington (2013) show that the primary determinant of the degree of alignment between the policy preferences of constituents and the legislative votes of their representatives is whether constituents are represented by a politician of their preferred or an opposing party. This implies that it is not their low socioeconomic status by itself that leads to a worse representation of

28. For this purpose, I define the political south as the 11 slave-holding states that declared secession from the US to form the Confederate States of America in the 1860s: Alabama, Arkansas, Florida, Georgia, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Texas, and Virginia.

29. These numbers are calculated based on item VCF0702 of the ANES cumulative file, using the provided ANES survey weights.

the financially disadvantaged, but rather their lower propensity to vote. Moreover, changes in the composition of the electorate due to the enfranchisement of specific population segments typically lead to large shifts in vote shares and policy outcomes (e.g., Miller, 2008; Cascio and Washington, 2014; Fujiwara, 2015).

Evidence from recent studies suggests that observed differences in electoral participation by socioeconomic status are at least partly driven by additional hurdles in the voting process for disadvantaged groups. For instance, Shue and Luttmer (2009) document that the prevalence of voting errors is higher in precincts where a higher fraction of the population lives in poverty or has low educational attainment, suggesting a higher cognitive cost of voting for people of low socioeconomic status. Similarly, Chen et al. (Forthcoming) show that voters in areas with a high proportion of Black or Hispanic residents are more likely to incur long wait times at polling places than residents of predominantly White neighborhoods. It is plausible that the same factors that cause observed turnout gaps by socioeconomic status also moderate the effect of experienced political competition on the propensity to vote.

To assess whether the magnitude of the effect of political experiences varies by socioeconomic status, I estimate variants of the estimation equation in Equation 1.1 that include an interaction term between experienced political competition and measures of education and income at the time of the survey. The results are presented in Table 1.2.

In columns (1) and (2), the interaction is with an indicator variable for holding a college degree and for belonging to the highest tercile of relative household income, respectively. The estimated coefficients indicate that experiencing low political competition reduces the voting propensity of respondents without a college education or with lower relative household income by about two percentage points. In contrast, the effect on more privileged respondents is statistically indistinguishable from zero (with F-test p-values of 0.602 and 0.661, respectively), though the interaction term is only significant for education. These differential effects explain roughly ten percent of the corresponding observed turnout gaps between respondents with high and low socioeconomic status in the ANES for the 2016 election.[30] Thus, a lack of political competition does not only contribute to voter apathy on average, but it also reinforces existing inequalities in electoral participation.

To assess whether the effect of political experiences fades with time, column (3) adds a linear trend for the number of years that have passed since age 25 at the time of the survey.[31] The negative effect of low experienced competition is estimated to be strongest for respondents just after their impressionable years, when they are 2.8 percentage points less likely to vote than respondents who experienced different parties win their state. Afterwards, each additional year

30. The observed turnout gaps are 15.7 percentage points between respondents in the highest tercile of household income and the remainder of the population and 19.5 percentage points between respondents with and without a college degree, respectively.

31. In unreported regressions, I implement a variant of the binning estimator proposed in Hainmueller, Mummolo, and Xu (2019) to gauge the validity of the functional form assumption of the linear interaction model. In the implementation, I use three tercile bins for the number of years passed (with respective bin medians of 7, 23 and 41). I fail to reject the null hypothesis that the additional parameters of the binned estimator are jointly equal to zero ($p = 0.300$), suggesting that the assumption of a linear interaction effect is a reasonable approximation within the range of common values of the number of years passed since age 25 in the data.

Table 1.2. Heterogeneity

	Dependent variable: Voted in last election		
	(1)	(2)	(3)
Experienced same state winner	−0.019***	−0.019**	−0.028**
	(0.007)	(0.008)	(0.011)
Experienced same state winner × College degree	0.025**		
	(0.012)		
Experienced same state winner × High income		0.015	
		(0.010)	
Experienced same state winner × Years passed			0.001*
			(0.000)
Demographic controls	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes
Observations	32 768	32 768	32 768
R^2 (within)	0.074	0.074	0.074
Mean (dependent variable)	0.714	0.714	0.714

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Years passed* are the number of years since age 25 at the time of the survey. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

diminishes the effect by approximately 0.1 percentage points on average. The estimated effect is still significantly different from zero 20 years after the impressionable years ($p = 0.019$ in an F-test). Thus, the linear approximation suggests that the effect of low experienced competition on individuals' propensity to vote persists for more than two decades before it has faded completely.

1.5.6 Additional Results

Additional results that are not immediately relevant to my hypothesis are presented in Appendix 1.B. I briefly summarize the findings from these analyses here.

Previous empirical research on the determinants of voter turnout has often found smaller effects for presidential elections (e.g., Gentzkow, 2006; Oberholzer-Gee and Waldfogel, 2009; Charles and Stephens Jr., 2013). To assess whether experienced political competition also has a differential effect on the propensity to vote in presidential and midterm elections, I conduct a sample split by election type, the results of which are reported in Table 1.B.2. I find that experiencing low political competition reduces turnout in midterm elections by four percentage points, while the effect for presidential elections is small and insignificant.[32] However, the interpretation of this result is plagued by differences in the methodology of the ANES Time Series Studies between the two election types: while only one interview is conducted with each respondent in the months after the election in midterm years, there is an additional pre-election interview in presidential years. Thus, it is unclear whether experienced political competition mostly affects elections with secondary electoral salience, or whether the differential effect is driven by an increase in social pressure to vote due to the pre-election interview before presidential elections (DellaVigna et al., 2017). By motivating counterfactual nonvoters to vote, this could reduce the scope for a manifestation of the effect of election experiences.[33]

To assess whether my results are biased towards zero due to measurement error in the state respondents lived in during their impressionable years, I also split the sample into movers and non-movers, with the latter defined as respondents for whom the state they grew up in equals their state of residence in the year of the interview. I would expect measurement error to play less of a role for non-movers, so a stronger effect for this group could indicate that my results should be interpreted as a lower bound on the true effect size. The results of this exercise, reported in Table 1.B.3, reveal that the estimated effect is indeed driven by non-movers. However, movers and non-movers probably also differ in other ways, and that could likewise explain the differential effect.

Following Giuliano and Spilimbergo (2014), I also estimate whether political experiences in other periods of life have an effect on subsequent voter turnout. Using the same methodology as in the main analysis, I construct indicator variables equal to one if the same party wins both elections in the state a respondent grew up in during age ranges 2 to 9, 10 to 17, 26 to 33, or 34 to 41. The coefficients of regressions of later-life turnout on these experiences, reported in Table 1.B.4, are all small and insignificant, in line with the idea that the impressionable years are the age span during which voting attitudes are shaped.[34]

In Table 1.B.5, I analyze whether experienced political competition affects party preferences, using both the self-reported party identification of all respondents and vote choice in the last presidential election conditional on casting a vote. For both outcomes, the coefficient on election experiences is small and insignificant, implying that voters who experience a lack of political competition during their impressionable years are not more or less likely to subsequently favor the Republican Party.

32. For midterm election turnout, I also run a regression with fixed effects for the current congressional district of residence interacted with election year instead of state interacted with election year to account for the fact that not all states have Senate races in a given midterm election. The results are very similar.

33. Jackman and Spahn (2018) analyze several potential causes for the overestimation of voter turnout in the ANES. They estimate that inadvertent mobilization due to the pre-election interview accounts for an increase in estimated turnout of at least three percentage points.

34. Since I only have data on the state respondents grew up in, measurement error is exacerbated the further away the calculated experiences are from adolescence. However, repeating the analysis for the subsample of non-movers does not change the interpretation of the results (see Table 1.C.5).

1.6 Mechanisms

The experiences of people who grow up under high and low political competition presumably differ in several ways, most notably a psychological component of feeling that their vote cannot meaningfully affect the election result, the experienced intensity of political campaigning, and the involvement of the electorate. Moreover, political experiences might affect individuals' subsequent voting propensity via different channels. In particular, economists and political scientists have discussed the role of political socialization, political identity, civic duty, habit formation, and beliefs about the benefit of voting as determinants of turnout, all of which could plausibly be affected by experienced electoral competition.

The aim of this section is to investigate the empirical validity of these channels. First, I analyze whether the effect of experienced political competition on subsequent voting inclination is reducible to differences in experienced involvement of the electorate. Second, I provide suggestive evidence on the causal mechanisms underlying the effect, within the limits of data availability in the ANES surveys. Since the strong assumptions for mediation analysis are unlikely to be satisfied in my setting (Huber, 2020), the general strategy is to test whether election experiences affect the potential mediating variables rather than adding them as independent variables to the main specification.

1.6.1 Is the Effect Driven by Experienced Political Participation?

Which facets of individuals' political competition experiences in particular leave a lasting imprint on their voting behavior? While historical data on the intensity of political campaigning by state is not easily available, I can proxy for the experienced involvement of the electorate using aggregate voter turnout. To this end, I gather data on the state-level aggregate turnout rate for all presidential elections between 1900 and 2009 from Gans (2011). As the denominator, this turnout rate uses the citizen-eligible population, i.e., the number of eligible voters is estimated from decennial census data taking into account state-level eligibility criteria based on gender, age, and race, and corrected for non-citizens residing in the US. While this adjustment is not perfect, it makes turnout rates as comparable as possible across states and election years for the given time period, alleviating concerns that differences are driven by changes in the share of the population that is ineligible to vote.[35]

The first two columns of Table 1.3 describe the relationship between the average state-level turnout rate experienced during the impressionable years and my measure of experienced political competition, once in a regression without any controls or fixed effects and once using the main specification. Not surprisingly, aggregate turnout is indeed about 1.5 to 2.1 percentage points lower if the same party wins both presidential elections in a given state.

If it was the experienced higher involvement of other eligible voters that affects later-life voting behavior, e.g., by setting a stronger social norm, one would expect the effect of expe-

35. Ideally, one would also like to exclude disenfranchised felons and individuals who have been declared mentally incompetent, and add eligible citizens from overseas, the share of which may vary across states and elections (McDonald and Popkin, 2001; McDonald, 2002). However, sufficient data for these adjustments is either not available on the state-level at all or only covers the recent past.

rienced political competition on the propensity to vote to shrink towards zero when experienced turnout is added to the regression as an additional independent variable. The results of this exercise are reported in column (3) of the table. The estimated coefficients indicate that an increase in experienced turnout of ten percentage points leads to a 0.5 percentage points higher likelihood of voting later in life ($p = 0.171$), while the coefficient on experienced political competition is similar to that in the main specification. Thus, there seems to be something in individuals' political experiences beyond the political involvement of other eligible voters that affects their subsequent voting behavior.

Table 1.3. Experienced Political Competition and Experienced Turnout

Dependent variable:	Experienced turnout rate		Voted in last election
	(1)	(2)	(3)
Experienced same state winner	−0.021*	−0.015***	−0.013**
	(0.012)	(0.005)	(0.006)
Experienced turnout rate			0.053
			(0.039)
Demographic controls	No	Yes	Yes
State × Year FE	No	Yes	Yes
Age FE	No	Yes	Yes
Birth cohort FE	No	Yes	Yes
State-when-young FE	No	Yes	Yes
Observations	32 569	32 550	32 550
R^2 (within)	0.005	0.019	0.073
Mean (dependent variable)	0.569	0.569	0.715

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced turnout rate* is the average turnout rate over both presidential elections in the state the respondent grew up in, coded as a fraction. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

1.6.2 Parental Voting Behavior

Regarding potential mechanisms underlying the effect of experienced competition on subsequent voting propensity, I first consider the possibility that the effect is driven by an intergenerational transmission of parental voting behavior or political knowledge. It is well known in political science that the family plays an important role for the political socialization of future generations (e.g., Verba, Schlozman, and Burns, 2005; Jennings, Stoker, and Bowers, 2009). In line with economic models of voting, parents who experience high political competition during their child's impressionable years might acquire more political information and be more likely to vote in these elections themselves, and then transmit their voting behavior and knowledge to the child. For instance, parents could directly teach their children about the importance of voting or act as role models when the whole family goes to the polls together.

To assess the plausibility of this channel, I make use of additional survey items from the ANES cumulative file. In a few survey years, the ANES includes information about the political interest of respondents' parents. I use standardized responses to questions VCF0308 and VCF0309, coded such that higher values imply higher political interest:

Father's political interest: *"Do you remember when you were growing up whether your father was very much interested in politics, somewhat interested, or didn't he pay much attention to it?"*

Mother's political interest: *"Now how about your mother? When you were growing up was she very much interested in politics, somewhat interested, or didn't she pay much attention to it?"*

With the caveat that these survey items are only available for less than 5000 respondents, they allow me to assess whether there is a relationship between individuals' political competition experiences during the impressionable years and their parents' interest in politics. To this end, I regress my measures of parental political interest on respondents' experienced political competition in columns (1) and (2) of Table 1.4, using the same controls and fixed effects as for the main results.

The estimated coefficients are close to zero, have mixed signs, and are far from statistically significant, implying that parental political interest is not lower on average for respondents who experience low political competition in their impressionable years.[36] Thus, the available data provides no indication that the effect of experienced political competition works through political socialization within the family.

1.6.3 Political Identity

The effect of experienced political competition on subsequent political participation could also be driven by an increase in the strength of individuals' political identity. When growing up in a highly competitive political atmosphere, young adults are more likely to encounter differing political views, both in their daily lives and due to political advertising and campaign activities of both parties during the campaign season. This might persistently strengthen the salience of individuals' political identity, which in turn makes it more likely that they vote in subsequent elections, e.g., because of an increase in expressive benefits from voting (e.g., Pons and Tricaud, 2018).

I measure respondents' strength of political identity using item VCF0305 from the ANES cumulative file. This item is collapsed from the self-placement of respondents on a 7-point party identification scale, where a value of 4 corresponds to a non-leaning independent, and values of 1 and 7 correspond to a strong identification with one of the two major parties. The responses are recoded such that a higher value implies a more extreme position on the scale, and I standardize the variable for ease of interpretation.

36. The results barely change if I only control for gender and minority status.

Table 1.4. Evidence on Mechanisms: Parental Voting Behavior, Political Identity, and Civic Duty

Dependent variable:	Father's political interest	Mother's political interest	Strength of party identification	Disagree: not important to vote
	(1)	(2)	(3)	(4)
Experienced same state winner	−0.005	0.006	0.008	−0.025
	(0.035)	(0.034)	(0.015)	(0.024)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes
Observations	4561	4663	32 593	7646
R^2 (within)	0.048	0.025	0.023	0.062

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. The dependent variables are standardized. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

In column (3) of Table 1.4, I report the coefficient estimate of a regression that resembles the main specification, but uses respondents' strength of political identity as the outcome variable. The results do not support the claim that experiencing more political competition during the impressionable years persistently increases the salience of individuals' political identity, so it seems unlikely that the effect on the propensity to vote is driven by this channel.

1.6.4 Civic Duty

One might reasonably expect the effect of experienced political competition on subsequent turnout to be mediated by a change in the perceived civic duty to vote. For instance, one interpretation of the enduring effects that are sometimes observed in academic get-out-the-vote field experiments is that the mobilization messages sent to the treatment groups persistently increase subjects' ethical concern for voting (Gerber and Green, 2017). Similarly, the higher political participation of other eligible voters might set a stronger descriptive voting norm. The voting norms taught in competitive political environments might also be more stringent than those that individuals learn in less competitive environments because of a higher perceived efficacy of voting.[37]

37. The opposite could also be true if voting norms and electoral competition are substitutes in promoting political participation.

As a measure of respondents' perceived civic duty to vote, I construct an indicator equal to one if they disagree with the following statement in the ANES:

> ***Not important to vote:*** *"Now I'd like to read some of the kinds of things people tell us when we interview them. Please tell me whether you agree or disagree with these statements: [...] It isn't so important to vote when you know your party doesn't have any chance to win." (item VCF0617)*

Column (4) of Table 1.4 reports the estimates of a regression of this indicator of perceived civic duty on election experiences, conditional on the fixed effects and controls of the main specification. For ease of comparison with the other investigated channels, the dependent variable is standardized.

The results are qualitatively in line with the proposed mechanism, but the coefficient is not significantly different from zero.[38] In light of the limited number of observations, I cannot dismiss the possibility that the effect of low experienced political competition on subsequent political participation is partly driven by a persistent decrease in the perceived civic duty to vote. However, the earlier finding that experienced political competition affects the likelihood of voting beyond experienced turnout suggests that a difference in descriptive voting norms is not the only mechanism at work.

1.6.5 Habit Formation

There is a large literature in political science that emphasizes the persistence of individual voting decisions over the life cycle and argues that turnout is habit-forming (e.g., Plutzer, 2002; Gerber, Green, and Shachar, 2003; Meredith, 2009). According to this view, turning out in one election increases the probability of voting in the next election.[39] In economics, Fujiwara, Meng, and Vogl (2016) provide evidence in support of habit formation in voter turnout using an exogenous decrease in the likelihood of voting due to rain on Election Day. In contrast, Bechtel, Hangartner, and Schmid (2018) find no indication of habit formation in the context of a compulsory voting intervention.

Habit formation is also a potential driver of the persistent effect of experienced political competition. According to economic models of voting, individuals who experience a highly competitive election when young should be more likely to vote already in this election. If turnout is habit-forming, this initial act of voting could then increase individuals' taste for voting in subsequent elections, and the resulting turnout gap between people with differing election experiences might endure for many years. This is the mental model put forward by Franklin (2004) as the underlying cause of long-term aggregate turnout dynamics. Unfortunately, the

38. The point estimate is equivalent to a 0.7 percentage points lower probability of disagreeing with the statement for people who experienced low political competition.

39. While a standard definition of habit formation in economics applied to voting would state that the act of voting persistently increases individuals' intrinsic taste for voting, Nickerson and Rogers (2014) report that political campaigns are also more likely to contact individuals who have registered to vote or voted in the past as part of their mobilization activities. This is because political campaigns rely on contact information from voter lists, which disproportionately include registered voters (Jackman and Spahn, 2021). Thus, habit formation in voting is often interpreted in a broad sense that encompasses the possibility of social interactions in existing studies, even though only the narrow definition fits the impressionable years hypothesis.

ANES does not contain information about respondents' individual turnout decisions during their impressionable years, so I am unable to investigate the plausibility of this channel empirically.

1.6.6 Beliefs about the Benefit of Voting

The last mechanism I consider is belief formation about the benefit of voting. According to economic models, in particular the belief that the own vote can influence the election outcome is a key driver of turnout decisions (e.g., Downs, 1957). Existing empirical research in economics on this topic has focused on the perceived competitiveness of the current election. The results of laboratory experiments that exogenously vary the anticipated closeness of small-scale participation games mostly support the theoretical prediction that turnout is increasing in pivotality beliefs (e.g., Levine and Palfrey, 2007; Faravelli, Kalayci, and Pimienta, 2020). However, these studies typically use neutral language to avoid cueing subjects with voting norms, so it is not obvious that their conclusions carry over to field settings. Indeed, Gerber, Hoffman, et al. (2020) find little evidence of a causal effect of pivotality beliefs on the intention or decision to vote in two field experiments, casting doubt on the importance of perceived competitiveness in the context of large real-world elections.[40]

The impressionable years hypothesis suggests that individuals' beliefs about the benefit of voting are shaped by the political experiences they make during early adulthood instead of or in addition to the expected closeness of the current election. People could either learn about the value of voting from personal observation of the degree of political competition in their state or be persuaded by the mobilization efforts and campaign activities of the competing political parties. If attitudes towards voting are sticky and do not change much after the impressionable years, they might persist and affect the propensity to vote in later-life elections.

To investigate this channel, I use questions about respondents' attitudes towards voting and the political process. The ideal question would ask whether a respondent believes that her own vote can affect the outcome of the current election, i.e., about an individual's personal efficacy rather than about government responsiveness to elections in general or about the electoral influence of certain groups or social classes. However, such a question does not exist in the ANES. Therefore, I make use of the four survey items which I have identified to be most closely related to voting attitudes:[41]

> ***People like me have no say:*** *"Now I'd like to read some of the kinds of things people tell us when we interview them. Please tell me whether you agree or disagree with these statements: [...] People like me don't have any say about what the government does." (item VCF0613)*
>
> ***Too many voters to care:*** *"Now I'd like to read some of the kinds of things people tell us when we interview them. Please tell me whether you agree or disagree with these statements: [...] So many*

40. Related field experiments in political science also deliver null results (e.g., Enos and Fowler, 2014). In contrast, Morton et al. (2015) and Bursztyn et al. (2020) use plausibly exogenous variation in exposure to pre-election or exit polls to provide observational evidence in favor of an effect of anticipated closeness on the decision to vote. However, it is difficult to fully rule out that their results are driven by reactions of political elites or changes in the perceived social pressure to vote rather than pivotality beliefs alone.

41. For some of the questions, there were minor changes in wording between survey years.

other people vote in the national elections that it doesn't matter much to me whether I vote or not." (item VCF0615)

Voting matters: *"Some people say that no matter who people vote for, it won't make any difference to what happens. Others say that who people vote for can make a big difference to what happens. Using the scale in the booklet, (where one means that voting won't make any difference to what happens and five means that voting can make a big difference), where would you place yourself?" (item VCF9250)*

Satisfied with democracy in US: *"On the whole, are you (very) satisfied, fairly satisfied, not very satisfied, or not at all satisfied with the way democracy works in the United States?" (items VCF9254 and VCF9255)*

For the first two questions, I have data on whether respondents agree or disagree with the given statement, while responses to the third and fourth question are on a 5-point and 4-point scale, respectively. I recode all variables such that higher values imply more favorable attitudes towards voting and standardize them for the sake of comparability.

The results of regressions of the main specification with the different belief outcomes as dependent variables are presented in the first four columns of Table 1.5. With the exception of the first outcome, for which the estimated effect is close to zero, the results are qualitatively in line with the mechanism. Experiencing low political competition during the impressionable years decreases later-life beliefs about the value of voting by between 1.8 and 4.7 percent of a standard deviation, holding constant the current electoral context. However, the effect is noisily estimated and only marginally significant for the one question that asks specifically whether voting can make a difference.

To alleviate concerns about multiple hypothesis testing, I also estimate the average effect size (AES) across all four belief outcomes, following Kling and Liebman (2004) and Clingingsmith, Khwaja, and Kremer (2009). For estimated election experiences coefficients β_j and respective standard deviations σ_j of the individual belief outcomes $Y_j, j \in \{1, 2, 3, 4\}$, the AES is equal to $\frac{1}{4}\sum_{j=1}^{4}\frac{\beta_j}{\sigma_j}$. In my case, σ_j equals one for all j because the belief outcomes are already standardized. To test against the null hypothesis of no average effect, I estimate the AES in a seemingly unrelated regression framework, which accounts for the covariance between the outcomes.[42] Compared to standard summary indices, the AES has the advantage that it does not require an overlapping sample, so it can be used even though only a subset of the four belief questions was asked in each survey year.

The estimated AES is reported in column (5) of the table. It supports the conclusion of the outcome-by-outcome regressions: experiencing low political competition persistently reduces beliefs about the benefit of voting, but the effect is noisily estimated due to the small number of observations. This finding is consistent with a belief-based explanation of the effect of political competition experiences on individuals' voting propensity in subsequent elections.

42. Specifically, I stack the four outcomes and then estimate a variant of Equation 1.1 in which each independent variable (including all fixed effects) is interacted with one indicator variable for each individual belief outcome. The resulting coefficients are the same as those in columns (1) to (4), but the regression additionally yields the necessary covariance matrix for hypothesis tests on the AES.

Table 1.5. Evidence on Mechanisms: Beliefs about the Benefit of Voting

Dependent variable:	Disagree: people like me have no say	Disagree: too many voters to care	Voting matters	Satisfied with democracy in US	AES (1) – (4)
	(1)	(2)	(3)	(4)	(5)
Experienced same state winner	0.002	−0.018	−0.047*	−0.036	−0.025*
	(0.016)	(0.025)	(0.028)	(0.028)	(0.014)
Demographic controls	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes	Yes
Observations	27 706	8826	8662	10 821	14 004
R^2 (within)	0.044	0.052	0.024	0.014	

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. The dependent variables in columns (1) to (4) are standardized. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. Column (5) reports the AES of *Experienced same state winner* across columns (1) to (4); the number of observations in this column refers to the average number of observations used in the estimation of the AES. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

1.7 Conclusion

In the context of the electoral system of the US, this chapter shows that individuals who experience a lack of political competition in their home state during their impressionable years are less likely to vote in subsequent elections. The effect is long-lasting and persists far into adulthood, indicating that periods of reduced political competition can contribute to voter apathy. In contrast to previous studies in political science, my fixed effects specification rules out that the documented effect is driven by the current electoral context, birth cohort effects, life cycle effects, or historical differences in voting behavior between states. Moreover, robustness of the effect to the inclusion of a large set of time-varying state-level controls and stability across different specifications further encourage a causal interpretation.

These results show that even variations in the political environment within an established democracy can have an enduring negative impact on the health of the political system. From a normative perspective, reductions in voter turnout are especially problematic if they lead to unrepresentative electoral participation (Lijphart, 1997). My findings indicate that the effect of experienced political competition on the propensity to vote is indeed heterogeneous, with a stronger impact on poorer and less educated individuals. This implies that a lack of political competition reinforces biases in electoral outcomes and resulting government policy towards citizens with high socioeconomic status, thus threatening the democratic ideal of equal representation. A related implication is that fluctuations in political competition over time can cause the policy preferences of particular cohorts to be underrepresented in the political system. Since there are substantial generational differences in vote choice (e.g., Ghitza, Gelman, and Auerbach, 2021), this bears the potential to induce additional distortions in political outcomes.

My findings are also relevant for economic models of voting, which tend to explain turnout with features of the current electoral context. My findings indicate that this focus is too narrow, as the propensity to vote is also shaped by personal experiences in the past. To understand long-term patterns in aggregate turnout, economic models would therefore have to incorporate historical factors, and the degree of political competition is one promising candidate.

Investigating the underlying mechanisms, I do not find any evidence that points to political socialization within the family or an increased salience of individuals' political identity as drivers of the effect. Instead, my findings are consistent with a persistent change in attitudes towards voting. Thus, one interpretation of the results is that some turnout-relevant beliefs about the benefit of voting are formed during the impressionable years rather than respond to the anticipated closeness of the current election.

Further research is needed to clarify how exactly experienced political competition affects later-life turnout. In particular, investigating the role of other plausible mechanisms like habit formation and quantifying their importance relative to persistent changes in voting attitudes would deepen our understanding of the determinants of voter apathy. In addition, it would be interesting to learn more about how exactly individuals' election experiences in competitive and uncompetitive states differ, and which factor is most important in determining subsequent voting inclination. One way to make progress in this regard is to find a source of exogenous variation in the potential mediating variables conditional on political competition, and then estimate the long-term effects on the propensity to vote.

Moreover, it remains an open question whether my findings are unique to the electoral system of the US or also generalize to other countries. One might expect fewer people to be deterred from voting due to election experiences in proportional electoral systems, where votes translate more directly into political representation and regional differences in perceived competitiveness are likely to be smaller.

Appendix 1.A Data Appendix

This appendix describes the source and construction of all variables that are not explained in detail in the main text.

1.A.1 Demographic Control Variables

In the main specification, I include a range of demographic control variables that are constructed from survey items in the ANES cumulative file. I control for gender by including a dummy variable equal to one for female respondents (based on item VCF0104), with male as the base category. To account for ethnicity, I include a minority dummy that is equal to one for respondents who are either Black, Hispanic, or belong to other or multiple races, such that White non-Hispanic is the omitted category (based on item VCF0105b). Education is captured by the inclusion of a dummy for completing at most high school and a dummy for completing college, with the base category consisting of respondents with below high school education (based on item VCF0110). Moreover, I add four dummies to control for respondents' place in the income distribution of the respective survey year (based on item VCF0114): one dummy for percentiles 17 to 33, one for percentiles 34 to 67, one for percentiles 68 to 95, and one for an income above the 95th percentile, with incomes below the 17th percentile as the base category. From item VCF0118, I construct a dummy equal to one for unemployed respondents, with employed, student, retired, and homemaker respondents as the omitted category. To account for religion, I include a dummy variable equal to one if the respondent considers herself to be Protestant, Catholic, or Jewish (based on item VCF0128), with no religion or smaller religious groups as the omitted category. Finally, I control for marital status by adding a dummy for married respondents (based on item VCF0147), with never married, divorced, separated, or widowed as the base category. For each category, I also include a dummy equal to one if data for that category is missing for the respondent to avoid losing observations due to missing control variables (as in Roth and Wohlfart, 2018).

1.A.2 Clustering to Account for Sampling Design

All reported standard errors are two-way cluster-robust, accounting for both correlated treatment assignment and complex survey sampling at the same time. To decrease sampling costs, the ANES employs a multi-stage sampling design in most election years. In a first stage, the US is partitioned into geographic regions called primary sampling units (PSUs), about 70 of which are randomly drawn to represent the whole population. In subsequent stages of the sampling process, individual respondents are drawn from within the sampled PSUs for face-to-face interviews. Thus, only a geographically clustered subset of the population is interviewed in each survey year. To account for the fact that some regions of the US do not show up in the sample, I group respondents into clusters according to the PSU they were sampled from.

If PSUs were sampled anew in each survey year, the individual ANES Time Series Studies would be independent, and the appropriate level of clustering would be by PSU interacted with survey year (Korn and Graubard, 1999). In reality, some ANES surveys were conducted with new respondents from previously drawn PSUs, such that two or more survey years rely on the same first-stage sampling frame. These respondents cannot be considered as belonging to different cluster groups. For the assignment of cluster groups, I therefore pool clusters from different survey years if they are from the same PSU and the surveys use the same sampling frame.

Since there is no data on the sampling design of the individual surveys in the ANES cumulative file, I construct this data based on information in the individual ANES Time Series Studies and the respective codebooks. In particular, the 1952, 1956, 1958, and 1960 ANES all use the same sampling design. The 1964 ANES is based on a new sampling frame that is reused for the 1966, 1968, and 1970 surveys. The study years 1972, 1974, and 1976 share the same sampling frame. A new sampling design is used for the 1978 and 1980 ANES, while the 1982 ANES uses an independent sampling design. The studies conducted in 1984, 1986, 1988, 1990, 1992, and 1994 share the same sampling design. A new frame is used for the studies conducted in 1996, 2000, and 2004.[43] The 2008, 2012, and 2016 ANES each use a new sampling design for the face-to-face interviews. In addition, the 2012 and 2016 studies also include an independent random sample of internet respondents. In the 2016 dataset, all eligible internet cases were randomly assigned to small groups, yielding clusters of between 8 and 20 respondents (The American National Election Studies, 2019a). To maintain consistency across survey years, I imitate this approach for the 2012 study and randomly combine the internet cases into 300 groups of size 12 or 13 instead of treating each respondent as a singleton cluster as coded in the original dataset (The American National Election Studies, 2016). In total, there are 1372 PSU clusters in the estimation sample.

1.A.3 Historical Control Variables

To assess the robustness of my main result, I make use of a large set of time-varying historical state characteristics that might be correlated with experienced political competition and could potentially affect the later-life turnout of residents. The data used to construct these variables are compiled from various sources.

To construct a time series of real per capita income growth for each state, I use annual estimates of per capita personal income at the state level since 1929 from U.S. Bureau of Economic Analysis (2019) and correct them for inflation by translating the figures into 2018 dollars based on Consumer Price Index data from U.S. Bureau of Labor Statistics (2019). State-level population estimates for the years 1900 to 2016 to calculate population growth are obtained from Watanabe (2017), who compiled the data from various publications of the US Census Bureau. To measure inequality, I use estimates of the share of fiscal income earned by the top five percent of the income distribution from the World Inequality Database (Alvaredo, Chancel, Piketty, Saez, and Zucman, 2019), which is available on the state level from 1918 to 2015. As a proxy for wealth, I gather data on home ownership rates since 1900 from the decennial census (U.S. Census Bureau, 2017), filling gaps via linear interpolation. I also use the IPUMS Current Population Survey (Flood, King, Rodgers, Ruggles, and Warren, 2018) for more fine-grained data

43. The 2000 study is special because it additionally utilizes an independent random sample of telephone interviews. I follow the approach in the dataset, which groups the telephone sample respondents into independent clusters based on their geographic location (The American National Election Studies, 2005a).

starting in 1977. Crime data on the state level are only available starting from 1960. I obtain estimates of both the property crime rate and the violent crime rate per 1000 inhabitants from U.S. Federal Bureau of Investigation Uniform Crime Reporting Statistics (2019).

For each respondent in the estimation sample, I calculate the average value of each of the state-level characteristics in the state they grew up in over their impressionable years. Additionally, I construct an indicator variable for experiencing a recession, following the approach of Giuliano and Spilimbergo (2014). This variable is equal to one if a respondent experienced at least one year with a per capita income growth of less than −3.4 percent in their home state during the impressionable years.

1.A.4 Other Outcomes

In Table 1.B.5, I analyze whether experienced political competition also affects party preferences, using two different outcome variables. I construct an indicator variable for identifying as a Republican from respondents' self-placement on a 7-point party identification scale (item VCF0301). The dummy is coded as one for people who think of themselves as weak or strong Republicans, as compared to independents (including those leaning to one of the two major parties) or weak or strong Democrats. I also use respondents' self-reported vote choice in the last presidential election (item VCF0705) to construct an indicator variable for voting Republican, as compared to voting Democrat or voting for a third party candidate, conditional on voting at all.

Appendix 1.B Additional Tables

Table 1.B.1. Descriptive Statistics for the Main Sample

Variable	Mean	SD	Min.	Median	Max.	Obs.
Voted in last election	0.726	0.446	0	1	1	32 787
Experienced same state winner	0.666	0.472	0	1	1	32 787
Experienced same national winner	0.585	0.493	0	1	1	32 787
Age	48.512	15.189	26	47	89	32 787
Birth cohort	1939.062	22.211	1879	1940	1990	32 787
Female	0.541	0.498	0	1	1	32 787
Gender missing	0.001	0.025	0	0	1	32 787
Minority	0.198	0.399	0	0	1	32 787
Race missing	0.003	0.054	0	0	1	32 787
High school	0.652	0.476	0	1	1	32 787
College degree	0.220	0.414	0	0	1	32 787
Education missing	0.005	0.074	0	0	1	32 787
Household income (percentile 17–33)	0.151	0.358	0	0	1	32 787
Household income (percentile 34–67)	0.313	0.464	0	0	1	32 787
Household income (percentile 68–95)	0.282	0.450	0	0	1	32 787
Household income (percentile 96–100)	0.051	0.219	0	0	1	32 787
Household income missing	0.060	0.238	0	0	1	32 787
Unemployed	0.079	0.269	0	0	1	32 787
Employment missing	0.031	0.172	0	0	1	32 787
Religious	0.858	0.349	0	1	1	32 787
Religion missing	0.003	0.058	0	0	1	32 787
Married	0.632	0.482	0	1	1	32 787
Marital status missing	0.007	0.083	0	0	1	32 787
Experienced turnout rate	0.568	0.148	0.074	0.600	0.932	32 569
Father’s political interest	2.135	0.799	1	2	3	4566
Mother’s political interest	1.687	0.764	1	2	3	4668
Strength of party identification	2.882	1.002	1	3	4	32 612
Disagree: not important to vote	0.908	0.289	0	1	1	7653
Disagree: people like me have no say	0.524	0.499	0	1	1	27 723
Disagree: too many voters to care	0.889	0.315	0	1	1	8838
Voting matters	4.020	1.133	1	4	5	8663
Satisfied with US democracy	2.839	0.771	1	3	4	10 831
Experienced same state winner (2–9 years)	0.673	0.469	0	1	1	32 276
Experienced same state winner (10–17 years)	0.661	0.473	0	1	1	32 622
Experienced same state winner (26–33 years)	0.677	0.467	0	1	1	32 569
Experienced same state winner (34–41 years)	0.686	0.464	0	1	1	31 376
Identify as Republican	0.259	0.438	0	0	1	32 415
Voted Republican in last presidential election	0.473	0.499	0	0	1	17 632
Experienced income growth	0.028	0.025	−0.026	0.022	0.177	28 835

(*continued on next page*)

Table 1.B.1. Descriptive Statistics for the Main Sample (*continued from previous page*)

Variable	Mean	SD	Min.	Median	Max.	Obs.
Experienced population growth	0.011	0.010	−0.027	0.010	0.082	32 735
Experienced inequality	0.256	0.045	0.124	0.244	0.535	31 928
Experienced a recession	0.429	0.495	0	0	1	28 835
Experienced wealth	0.582	0.108	0.287	0.598	0.788	32 746
Experienced property crime rate	39.663	13.759	6.966	39.818	86.947	15 566
Experienced violent crime rate	4.537	2.545	0.225	4.194	25.752	15 566

Notes: Unweighted mean, standard deviation (SD), minimum, median, maximum, and number of observations for each variable in the main estimation sample.

Table 1.B.2. Subsample Analyses for Turnout in Presidential and Midterm Elections

	Dependent variable: Voted in last election					
	Presidential elections subsample		Midterm elections subsample			
	(1)	(2)	(3)	(4)	(5)	(6)
Experienced same state winner	−0.003		−0.040***	−0.045***		
	(0.007)		(0.012)	(0.014)		
Experienced same national winner		−0.008			0.000	0.002
		(0.006)			(0.010)	(0.011)
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	No	Yes	No
Congressional district × Year FE	No	No	No	Yes	No	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	No	Yes	Yes	No	No
Cohort group FE	No	Yes	No	No	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	23 006	23 006	9761	9598	9761	9598
R^2 (within)	0.075	0.077	0.075	0.070	0.075	0.069
Mean (dependent variable)	0.776	0.776	0.578	0.579	0.578	0.579

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. Columns (1) and (2) use the subsample of ANES studies conducted in presidential election years, whereas columns (3) to (6) use the subsample of studies conducted in midterm election years. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced same national winner* is an indicator equal to one if the presidency was won by the same major party in both presidential elections experienced during the impressionable years. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.B.3. Subsample Analyses for Non-movers and Movers

	Dependent variable: Voted in last election			
	Non-movers subsample		Movers subsample	
	(1)	(2)	(3)	(4)
Experienced same state winner	−0.018***		0.003	
	(0.007)		(0.012)	
Experienced same national winner		−0.007		−0.002
		(0.006)		(0.011)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	No	Yes	No
Cohort group FE	No	Yes	No	Yes
State-when-young FE	No	No	Yes	Yes
Observations	22 893	22 893	9797	9798
R^2 (within)	0.078	0.078	0.072	0.071
Mean (dependent variable)	0.709	0.709	0.727	0.727

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. Columns (1) and (2) use the subsample of respondents for whom the state of residence at the time of the interview equals the state they grew up in, whereas columns (3) and (4) use the subsample of respondents who resided in a different state. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced same national winner* is an indicator equal to one if the presidency was won by the same major party in both presidential elections experienced during the impressionable years. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.B.4. Experiences in Different Periods of Life

	Dependent variable: Voted in last election			
	(1)	(2)	(3)	(4)
Experienced same state winner (2–9 years)	−0.006			
	(0.006)			
Experienced same state winner (10–17 years)		0.005		
		(0.006)		
Experienced same state winner (26–33 years)			0.000	
			(0.007)	
Experienced same state winner (34–41 years)				−0.001
				(0.008)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes
Observations	37 090	37 422	26 155	19 899
R^2 (within)	0.067	0.067	0.077	0.078
Mean (dependent variable)	0.678	0.679	0.746	0.763

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during the specified age range, indicating low political competition. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.B.5. Other Outcomes

Dependent variable:	Identify as Republican	Voted Republican in last presidential election	
	(1)	(2)	(3)
Experienced same state winner	0.005	−0.003	−0.007
	(0.007)	(0.009)	(0.007)
Demographic controls	Yes	Yes	Yes
Party identification dummies	No	No	Yes
State × Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes
Observations	32 394	17 614	17 614
R^2 (within)	0.049	0.106	0.410
Mean (dependent variable)	0.267	0.495	0.495

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. Columns (2) and (3) are restricted to the subsample of respondents who reported voting from ANES studies conducted in presidential election years. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Appendix 1.C Robustness

Table 1.C.1. Alternative Fixed Effects Specifications

	Dependent variable: Voted in last election			
	(1)	(2)	(3)	(4)
Experienced same state winner	−0.013**	−0.022***	−0.020***	−0.016**
	(0.006)	(0.007)	(0.007)	(0.007)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	No
Congressional district × Year FE	No	No	No	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	No	No	Yes
State-when-young FE	No	Yes	No	Yes
State-when-young × Cohort group FE	Yes	No	Yes	No
Region-when-young × Birth cohort FE	No	Yes	Yes	No
Observations	32 765	32 762	32 759	31 005
R^2 (within)	0.073	0.073	0.073	0.070
Mean (dependent variable)	0.714	0.714	0.714	0.713

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

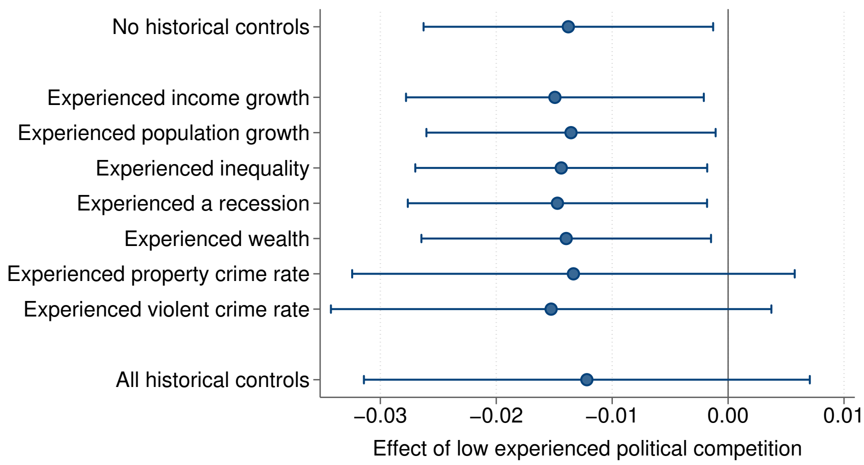

Figure 1.C.1. Effect of Low Experienced Political Competition with Historical Control Variables

Notes: Plot of coefficient estimates of the effect of experiencing low political competition during one's impressionable years on later-life voter turnout with different historical control variables. The estimates are from OLS regressions, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, of an indicator for having voted in the last election on *Experienced same state winner*, demographic control variables, state-year FE, age FE, birth cohort FE, state-when-young FE and the specified historical control variable(s). Error bars indicate 95 percent confidence intervals, constructed using two-way cluster-robust standard errors by state-when-young × election cohort and PSU. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced income growth* and *Experienced population growth* are average yearly growth rates of real per capita personal income and population, respectively. *Experienced inequality* is the average income share of the top five percent of the income distribution. *Experienced a recession* is an indicator equal to one if the respondent experienced at least one year with a real per capita personal income growth below −3.4 percent. *Experienced wealth* is the average home ownership rate. Experienced crime rates are per 1000 inhabitants. All experiences refer to the state the respondent grew up in during the period of their impressionable years. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies.

Table 1.C.2. Alternative Measures of Experienced Political Competition

	Dependent variable: Voted in last election			
	(1)	(2)	(3)	(4)
Experienced same state winner (incl. third parties)	−0.015**			
	(0.006)			
Experienced state victory margin (Q2)		−0.012		
		(0.008)		
Experienced state victory margin (Q3)		−0.016**		
		(0.008)		
Experienced state victory margin (Q4)		−0.012		
		(0.009)		
Experienced electoral victory margin (Q2)			−0.016**	
			(0.008)	
Experienced electoral victory margin (Q3)			−0.016**	
			(0.006)	
Experienced electoral victory margin (Q4)			−0.006	
			(0.007)	
Experienced national victory margin (Q2)				−0.007
				(0.008)
Experienced national victory margin (Q3)				0.000
				(0.010)
Experienced national victory margin (Q4)				−0.004
				(0.008)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	No	No
Cohort group FE	No	No	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes
Observations	32 768	32 768	32 768	32 768
R^2 (within)	0.073	0.073	0.075	0.074
Mean (dependent variable)	0.714	0.714	0.714	0.714

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. *Experienced same state winner (incl. third parties)* is an indicator equal to one if the state the respondent grew up in was won by the same party in both presidential elections experienced during their impressionable years, indicating low political competition. The coefficients in columns (2) to (4) indicate quartiles of the distribution of the respective measure of political competition in the estimation sample. *Experienced state victory margin* is the average popular vote victory margin in the state the respondent grew up in over the two presidential elections experienced during their impressionable years. *Experienced electoral victory margin* and *Experienced national victory margin* are the average margins of victory in electoral college vote and national popular vote over the two presidential elections experienced by the respondent during their impressionable years, respectively. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.C.3. Alternative Clustering of Standard Errors

	Dependent variable: Voted in last election			
	SE clustered for correlated treatment assignment only		SE clustered for complex survey sampling only	
	(1)	(2)	(3)	(4)
Experienced same state winner	−0.014**		−0.014**	
	(0.006)		(0.007)	
Experienced same national winner		−0.005		−0.005
		(0.005)		(0.005)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	No	Yes	No
Cohort group FE	No	Yes	No	Yes
State-when-young FE	Yes	Yes	Yes	Yes
Observations	32 768	32 768	32 768	32 768
R^2 (within)	0.073	0.074	0.073	0.074
Mean (dependent variable)	0.714	0.714	0.714	0.714

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with cluster-robust standard errors in parentheses. Standard errors are clustered by state-when-young × election cohort in columns (1) and (2), and by PSU in columns (3) and (4). *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. *Experienced same national winner* is an indicator equal to one if the presidency was won by the same major party in both presidential elections experienced during the impressionable years. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.C.4. Other Robustness Checks

	Dependent variable: Voted in last election			
	Full sample		Subsamples without political south	
	(1)	(2)	(3)	(4)
Experienced same state winner	−0.016**	−0.014**	−0.015*	−0.014*
	(0.006)	(0.006)	(0.008)	(0.008)
Demographic controls	Yes	Yes	Yes	Yes
Party identification dummies	No	Yes	No	No
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes	Yes
State-when-young FE	Yes	Yes	Yes	Yes
Observations	32 768	32 768	23 416	22 918
R^2 (within)	0.003	0.090	0.065	0.064
Mean (dependent variable)	0.714	0.714	0.749	0.747

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. The sample is restricted to respondents who did not grow up in states that belong to the political south of the US in column (3), and to respondents who did not reside in states that belong to the political south at the time of the survey in column (4). *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during their impressionable years, indicating low political competition. Demographic controls include only gender and minority status dummies in column (1), and additionally education, religiosity, marital status, and income group dummies in columns (2) to (4). * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 1.C.5. Experiences in Different Periods of Life (Non-movers Only)

	Dependent variable: Voted in last election			
	(1)	(2)	(3)	(4)
Experienced same state winner (2–9 years)	−0.005 (0.007)			
Experienced same state winner (10–17 years)		0.003 (0.006)		
Experienced same state winner (26–33 years)			0.003 (0.009)	
Experienced same state winner (34–41 years)				0.009 (0.009)
Demographic controls	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Birth cohort FE	Yes	Yes	Yes	Yes
Observations	27394	27607	17963	13515
R^2 (within)	0.069	0.069	0.083	0.083
Mean (dependent variable)	0.661	0.663	0.742	0.758

Notes: OLS estimates, weighted to be nationally representative of the US voting-eligible population in each election year using ANES survey weights, with two-way cluster-robust standard errors by state-when-young × election cohort and PSU in parentheses. The sample is restricted to respondents for whom the state of residence at the time of the interview equals the state they grew up in. *Experienced same state winner* is an indicator equal to one if the state the respondent grew up in was won by the same major party in both presidential elections experienced during the specified age range, indicating low political competition. Demographic controls include gender, minority status, education, religiosity, marital status, and income group dummies. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey Wooldridge. 2017. "When Should You Adjust Standard Errors for Clustering?" NBER Working Paper No. 24003. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w24003. [17]

Acemoglu, Daron, Nicolás Ajzenman, Cevat Giray Aksoy, Martin Fiszbein, and Carlos A. Molina. 2021. "(Successful) Democracies Breed Their Own Support." NBER Working Paper No. 29167. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w29167. [8]

Akbulut-Yuksel, Mevlude, Dozie Okoye, and Mutlu Yuksel. 2020. "Social Changes in Impressionable Years and Adult Political Attitudes: Evidence from Jewish Expulsions in Nazi Germany." *Economic Inquiry* 58 (1): 184–208. DOI: 10.1111/ecin.12843. [8, 21]

Alacevich, Caterina, and Dijana Zejcirovic. 2020. "Does Violence against Civilians Depress Voter Turnout? Evidence from Bosnia and Herzegovina." *Journal of Comparative Economics* 48 (4): 841–65. DOI: 10.1016/j.jce.2020.04.006. [8]

Alesina, Alberto, and Nicola Fuchs-Schündeln. 2007. "Good-Bye Lenin (or Not?): The Effect of Communism on People's Preferences." *American Economic Review* 97 (4): 1507–28. DOI: 10.1257/aer.97.4.1507. [8]

Ali, S. Nageeb, and Charles Lin. 2013. "Why People Vote: Ethical Motives and Social Incentives." *American Economic Journal: Microeconomics* 5 (2): 73–98. DOI: 10.1257/mic.5.2.73. [10]

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. 2019. "World Inequality Database [online]." Available at https://wid.world/data/. Accessed December 4, 2019. [36]

Ashenfelter, Orley, and Stanley Kelley Jr. 1975. "Determinants of Participation in Presidential Elections." *Journal of Law and Economics* 18 (3): 695–733. DOI: 10.1086/466834. [15]

Bafumi, Joseph, and Andrew Gelman. 2006. "Fitting Multilevel Models When Predictors and Group Effects Correlate." Unpublished manuscript, Columbia University. URL: http://www.stat.columbia.edu/~gelman/research/unpublished/Bafumi_Gelman_Midwest06.pdf. [9]

Balles, Patrick, Ulrich Matter, and Alois Stutzer. 2022. "Television Market Size and Political Accountability in the US House of Representatives." IZA Discussion Paper No. 15277. Bonn: IZA – Institute of Labor Economics. URL: https://docs.iza.org/dp15277.pdf. [7]

Bechtel, Michael M., Dominik Hangartner, and Lukas Schmid. 2018. "Compulsory Voting, Habit Formation, and Political Participation." *Review of Economics and Statistics* 100 (3): 467–76. DOI: 10.1162/rest_a_00701. [8, 30]

Becker, Johannes, Andreas Peichl, and Johannes Rincke. 2009. "Politicians' Outside Earnings and Electoral Competition." *Public Choice* 140 (3–4): 379–94. DOI: 10.1007/s11127-009-9426-y. [7]

Bell, Andrew, and Kelvyn Jones. 2015. "Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data." *Political Science Research and Methods* 3 (1): 133–53. DOI: 10.1017/psrm.2014.7. [9]

Besley, Timothy, Torsten Persson, and Daniel M. Sturm. 2010. "Political Competition, Policy and Growth: Theory and Evidence from the US." *Review of Economic Studies* 77 (4): 1329–52. DOI: 10.1111/j.1467-937X.2010.00606.x. [7, 20, 22]

Blais, André, and Daniel Rubenson. 2013. "The Source of Turnout Decline: New Values or New Contexts?" *Comparative Political Studies* 46 (1): 95–117. DOI: 10.1177/0010414012453032. [9]

Blanco, Luisa, and Isabel Ruiz. 2013. "The Impact of Crime and Insecurity on Trust in Democracy and Institutions." *American Economic Review: Papers & Proceedings* 103 (3): 284–88. DOI: 10.1257/aer.103.3.284. [20]

Brunner, Eric, Stephen L. Ross, and Ebonya Washington. 2013. "Does Less Income Mean Less Representation?" *American Economic Journal: Economic Policy* 5 (2): 53–76. DOI: 10.1257/pol.5.2.53. [5, 22]

Buonanno, Paolo, Daniel Montolio, and Paolo Vanin. 2009. "Does Social Capital Reduce Crime?" *Journal of Law and Economics* 52 (1): 145–70. DOI: 10.1086/595698. [20]

Bursztyn, Leonardo, Davide Cantoni, Patricia Funk, Felix Schönenberger, and Noam Yuchtman. 2020. "Identifying the Effect of Election Closeness on Voter Turnout: Evidence from Swiss Referenda." NBER Working Paper No. 23490. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w23490. [9, 31]

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2011. "Robust Inference with Multiway Clustering." *Journal of Business & Economic Statistics* 29 (2): 2011. DOI: 10.1198/jbes.2010.07136. [17]

Cantoni, Enrico. 2020. "A Precinct Too Far: Turnout and Voting Costs." *American Economic Journal: Applied Economics* 12 (1): 61–85. DOI: 10.1257/app.20180306. [8, 18]

Cantoni, Enrico, and Vincent Pons. 2021. "Strict ID Laws Don't Stop Voters: Evidence from a U.S. Nationwide Panel, 2008–2018." *Quarterly Journal of Economics* 136 (4): 2615–60. DOI: 10.1093/qje/qjab019. [8]

Cantoni, Enrico, and Vincent Pons. 2022. "Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the United States." *American Economic Review* 112 (4): 1226–72. DOI: 10.1257/aer.20201660. [16]

Carreri, Maria, and Edoardo Teso. Forthcoming. "Economic Recessions and Congressional Preferences for Redistribution." *Review of Economics and Statistics*, DOI: 10.1162/rest_a_01053. [8]

Cascio, Elizabeth U., and Ebonya Washington. 2014. "Valuing the Vote: The Redistribution of Voting Rights and State Funds following the Voting Rights Act of 1965." *Quarterly Journal of Economics* 129 (1): 379–433. DOI: 10.1093/qje/qjt028. [23]

Charles, Kerwin Kofi, and Melvin Stephens Jr. 2013. "Employment, Wages, and Voter Turnout." *American Economic Journal: Applied Economics* 5 (4): 111–43. DOI: 10.1257/app.5.4.111. [8, 20, 24]

Chen, M. Keith, Kareem Haggag, Devin G. Pope, and Ryne Rohla. Forthcoming. "Racial Disparities in Voting Wait Times: Evidence from Smartphone Data." *Review of Economics and Statistics*, DOI: 10.1162/rest_a_01012. [8, 23]

Chyn, Eric, and Kareem Haggag. 2019. "Moved to Vote: The Long-Run Effects of Neighborhoods on Political Participation." NBER Working Paper No. 26515. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w26515. [8, 18]

Clingingsmith, David, Asim Ijaz Khwaja, and Michael Kremer. 2009. "Estimating the Impact of the Hajj: Religion and Tolerance in Islam's Global Gathering." *Quarterly Journal of Economics* 124 (3): 1133–70. DOI: 10.1162/qjec.2009.124.3.1133. [32]

Cohodes, Sarah, and James J. Feigenbaum. 2021. "Why Does Education Increase Voting? Evidence from Boston's Charter Schools." NBER Working Paper No. 29308. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w29308. [8]

Correia, Sergio. 2016. "A Feasible Estimator for Linear Models with Multi-Way Fixed Effects." Unpublished manuscript, Duke University. March 2016. URL: http://scorreia.com/research/hdfe.pdf. [15]

Dal Bó, Ernesto, Pedro Dal Bó, and Jason Snyder. 2009. "Political Dynasties." *Review of Economic Studies* 76 (1): 115–42. DOI: 10.1111/j.1467-937X.2008.00519.x. [7]

De Paola, Maria, and Vincenzo Scoppa. 2011. "Political Competition and Politician Quality: Evidence from Italian Municipalities." *Public Choice* 148 (3–4): 547–59. DOI: 10.1007/s11127-010-9683-9. [7]

DellaVigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2017. "Voting to Tell Others." *Review of Economic Studies* 84 (1): 143–81. DOI: 10.1093/restud/rdw056. [8, 11, 25]

Di Tella, Rafael, Sebastian Galiani, and Ernesto Schargrodsky. 2007. "The Formation of Beliefs: Evidence from the Allocation of Land Titles to Squatters." *Quarterly Journal of Economics* 122 (1): 209–41. DOI: 10.1162/qjec.122.1.209. [8]

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York, NY: Harper & Row. [10, 31]

Eichengreen, Barry, Orkun Saka, and Cevat Giray Aksoy. 2022. "The Political Scar of Epidemics." NBER Working Paper No. 27401. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w27401. [8]

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. [8]

Enos, Ryan D., and Anthony Fowler. 2014. "Pivotality and Turnout: Evidence from a Field Experiment in the Aftermath of a Tied Election." *Political Science Research and Methods* 2 (2): 309–19. DOI: 10.1017/psrm.2014.5. [31]

Falck, Oliver, Robert Gold, and Stephan Heblich. 2014. "E-lections: Voting Behavior and the Internet." *American Economic Review* 104 (7): 2238–65. DOI: 10.1257/aer.104.7.2238. [8]

Faravelli, Marco, Kenan Kalayci, and Carlos Pimienta. 2020. "Costly Voting: A Large-Scale Real Effort Experiment." *Experimental Economics* 23 (2): 468–92. DOI: 10.1007/s10683-019-09620-3. [8, 31]

Feddersen, Timothy, and Alvaro Sandroni. 2006. "A Theory of Participation in Elections." *American Economic Review* 96 (4): 1271–82. DOI: 10.1257/aer.96.4.1271. [10]

Feddersen, Timothy J. 2004. "Rational Choice Theory and the Paradox of Not Voting." *Journal of Economic Perspectives* 18 (1): 99–112. DOI: 10.1257/089533004773563458. [5]

Ferraz, Claudio, and Frederico Finan. 2011. "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments." *American Economic Review* 101 (4): 1274–311. DOI: 10.1257/aer.101.4.1274. [7]

Fleck, Robert K. 1999. "The Value of the Vote: A Model and Test of the Effects of Turnout on Distributive Policy." *Economic Inquiry* 37 (4): 609–23. DOI: 10.1111/j.1465-7295.1999.tb01451.x. [5]

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2018. "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]." Minneapolis, MN: IPUMS. DOI: 10.18128/D030.V6.0. [36]

Franklin, Mark N. 2004. *Voter Turnout and the Dynamics of Electoral Competition in Established Democracies since 1945*. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9780511616884. [9, 30]

Franklin, Mark N., Patrick Lyons, and Michael Marsh. 2004. "Generational Basis of Turnout Decline in Established Democracies." *Acta Politica* 39 (2): 115–51. DOI: 10.1057/palgrave.ap.5500060. [9]

Fuchs-Schündeln, Nicola, and Matthias Schündeln. 2015. "On the Endogeneity of Political Preferences: Evidence from Individual Experience with Democracy." *Science* 347 (6226): 1145–48. DOI: 10.1126/science.aaa0880. [8]

Fuhrmann, Delia, Lisa J. Knoll, and Sarah-Jayne Blakemore. 2015. "Adolescence as a Sensitive Period of Brain Development." *Trends in Cognitive Sciences* 19 (10): 558–66. DOI: 10.1016/j.tics.2015.07.008. [10]

Fujiwara, Thomas. 2015. "Voting Technology, Political Responsiveness, and Infant Health: Evidence from Brazil." *Econometrica* 83 (2): 423–64. DOI: 10.3982/ECTA11520. [23]

Fujiwara, Thomas, Kyle Meng, and Tom Vogl. 2016. "Habit Formation in Voting: Evidence from Rainy Elections." *American Economic Journal: Applied Economics* 8 (4): 160–88. DOI: 10.1257/app.20140533. [8, 30]

Gans, Curtis, editor. 2011. *Voter Turnout in the United States 1788–2009*. Washington, DC: CQ Press. DOI: 10.4135/9781608712700. [26]

Gentzkow, Matthew. 2006. "Television and Voter Turnout." *Quarterly Journal of Economics* 121 (3): 931–72. DOI: 10.1162/qjec.121.3.931. [8, 24]

Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson. 2011. "The Effect of Newspaper Entry and Exit on Electoral Politics." *American Economic Review* 101 (7): 2980–3018. DOI: 10.1257/aer.101.7.2980. [8]

Gerber, Alan, Mitchell Hoffman, John Morgan, and Collin Raymond. 2020. "One in a Million: Field Experiments on Perceived Closeness of the Election and Voter Turnout." *American Economic Journal: Applied Economics* 12 (3): 287–325. DOI: 10.1257/app.20180574. [8, 31]

Gerber, Alan S., and Donald P. Green. 2017. "Field Experiments on Voter Mobilization: An Overview of a Burgeoning Literature." In *Handbook of Economic Field Experiments*. Edited by Abhijit V. Banerjee and Esther Duflo. Vol. 1, Amsterdam: North-Holland. Chapter 9, 395–438. DOI: 10.1016/bs.hefe.2016.09.002. [29]

Gerber, Alan S., Donald P. Green, and Ron Shachar. 2003. "Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment." *American Journal of Political Science* 47 (3): 540–50. DOI: 10.1111/1540-5907.00038. [30]

Ghitza, Yair, Andrew Gelman, and Jonathan Auerbach. 2021. “The Great Society, Reagan’s Revolution, and Generations of Presidential Voting.” Unpublished manuscript, Columbia University. December 9, 2021. URL: http://www.stat.columbia.edu/~gelman/research/published/cohort_voting_20211209.pdf. [10, 34]

Giuliano, Paola, and Antonio Spilimbergo. 2014. “Growing Up in a Recession.” *Review of Economic Studies* 81 (2): 787–817. DOI: 10.1093/restud/rdt040. [8, 10, 15, 25, 37]

Glaeser, Edward L., Giacomo A. M. Ponzetto, and Jesse M. Shapiro. 2005. “Strategic Extremism: Why Republicans and Democrats Divide on Religious Values.” *Quarterly Journal of Economics* 120 (4): 1283–330. DOI: 10.1162/003355305775097533. [15]

Górecki, Maciej A. 2013. “Election Closeness, Habit Formation and Voter Turnout: Evidence from Sixteen Swedish Elections.” *Political Studies* 61 (S1): 234–48. DOI: 10.1111/1467-9248.12017. [9]

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. DOI: 10.1017/pan.2018.46. [23]

Heath, Oliver. 2007. “Explaining Turnout Decline in Britain, 1964–2005: Party Identification and the Political Context.” *Political Behavior* 29 (4): 493–516. DOI: 10.1007/s11109-007-9039-4. [9]

Hodler, Roland, Simon Luechinger, and Alois Stutzer. 2015. “The Effects of Voting Costs on the Democratic Process and Public Finances.” *American Economic Journal: Economic Policy* 7 (1): 141–71. DOI: 10.1257/pol.20120383. [8]

Huber, Martin. 2020. “Mediation Analysis.” In *Handbook of Labor, Human Resources and Population Economics*. Edited by Klaus F. Zimmermann. Cham: Springer. DOI: 10.1007/978-3-319-57365-6_162-2. [26]

Jackman, Simon, and Bradley Spahn. 2018. “Why Does the American National Election Study Overestimate Voter Turnout?” *Political Analysis* 27 (2): 193–207. DOI: 10.1017/pan.2018.36. [25]

Jackman, Simon, and Bradley Spahn. 2021. “Politically Invisible in America.” *PS: Political Science & Politics* 54 (4): 623–29. DOI: 10.1017/S1049096521000639. [11, 30]

Jennings, M. Kent, Laura Stoker, and Jake Bowers. 2009. “Politics across Generations: Family Transmission Reexamined.” *Journal of Politics* 71 (3): 782–99. DOI: 10.1017/S0022381609090719. [27]

Kaplan, Ethan, and Haishan Yuan. 2020. “Early Voting Laws, Voter Turnout, and Partisan Vote Composition: Evidence from Ohio.” *American Economic Journal: Applied Economics* 12 (1): 32–60. DOI: 10.1257/app.20180192. [8, 18]

Kling, Jeffrey R., and Jeffrey B. Liebman. 2004. “Experimental Analysis of Neighborhood Effects on Youth.” IRS Working Papers No. 483. Princeton, NJ: Industrial Relations Section, Princeton University. URL: https://arks.princeton.edu/ark:/88435/dsp01m613mx58m. [32]

Korn, Edward L., and Barry I. Graubard. 1999. *Analysis of Health Surveys*. New York, NY: John Wiley & Sons. DOI: 10.1002/9781118032619. [17, 35]

Krosnick, Jon A., and Duane F. Alwin. 1989. “Aging and Susceptibility to Attitude Change.” *Journal of Personality and Social Psychology* 57 (3): 416–25. DOI: 10.1037/0022-3514.57.3.416. [6, 10]

Ledyard, John O. 1984. “The Pure Theory of Large Two-Candidate Elections.” *Public Choice* 44 (1): 7–41. DOI: 10.1007/BF00124816. [10]

Leighley, Jan E., and Jonathan Nagler. 2013. *Who Votes Now?: Demographics, Issues, Inequality, and Turnout in the United States*. Princeton, NJ: Princeton University Press. DOI: 10.23943/princeton/9780691159348.001.0001. [15]

Levine, David K., and Thomas R. Palfrey. 2007. “The Paradox of Voter Participation? A Laboratory Study.” *American Political Science Review* 101 (1): 143–58. DOI: 10.1017.S0003055407070013. [8, 31]

Lijphart, Arend. 1997. “Unequal Participation: Democracy’s Unresolved Dilemma.” *American Political Science Review* 91 (1): 1–14. DOI: 10.2307/2952255. [5, 22, 33]

Lott Jr., John R., and Lawrence W. Kenny. 1999. “Did Women’s Suffrage Change the Size and Scope of Government?” *Journal of Political Economy* 107 (6): 1163–98. DOI: 10.1086/250093. [13]

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics* 126 (1): 373–416. DOI: 10.1093/qje/qjq004. [8]

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics* 131 (1): 53–87. DOI: 10.1093/qje/qjv037. [8]

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. 2021. "The Making of Hawks and Doves." *Journal of Monetary Economics* 117: 19–42. DOI: 10.1016/j.jmoneco.2020.04.002. [8]

Malmendier, Ulrike, and Leslie Sheng Shen. 2018. "Scarred Consumption." NBER Working Paper No. 24696. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w24696. [8]

McDonald, Michael P. 2002. "The Turnout Rate among Eligible Voters in the States, 1980–2000." *State Politics & Policy Quarterly* 2 (2): 199–212. DOI: 10.1177/153244000200200205. [26]

McDonald, Michael P., and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95 (4): 963–74. DOI: 10.1017/S0003055400400134. [26]

Meredith, Marc. 2009. "Persistence in Political Participation." *Quarterly Journal of Political Science* 4 (3): 187–209. DOI: 10.1561/100.00009015. [30]

Miller, Grant. 2008. "Women's Suffrage, Political Responsiveness, and Child Survival in American History." *Quarterly Journal of Economics* 123 (3): 1287–327. DOI: 10.1162/qjec.2008.123.3.1287. [23]

Morton, Rebecca B., Daniel Muller, Lionel Page, and Benno Torgler. 2015. "Exit Polls, Turnout, and Bandwagon Voting: Evidence from a Natural Experiment." *European Economic Review* 77: 65–81. DOI: 10.1016/j.euroecorev.2015.03.012. [9, 31]

Mueller, Dennis C., and Thomas Stratmann. 2003. "The Economic Effects of Democratic Participation." *Journal of Public Economics* 87 (9–10): 2129–55. DOI: 10.1016/S0047-2727(02)00046-4. [5]

Naidu, Suresh. 2012. "Suffrage, Schooling, and Sorting in the Post-Bellum U.S. South." NBER Working Paper No. 18129. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w18129. [22]

Nickerson, David W., and Todd Rogers. 2014. "Political Campaigns and Big Data." *Journal of Economic Perspectives* 28 (2): 51–74. DOI: 10.1257/jep.28.2.51. [30]

Oberholzer-Gee, Felix, and Joel Waldfogel. 2009. "Media Markets and Localism: Does Local News en Español Boost Hispanic Voter Turnout?" *American Economic Review* 99 (5): 2120–28. DOI: 10.1257/aer.99.5.2120. [24]

Pacheco, Julianna Sandell. 2008. "Political Socialization in Context: The Effect of Political Competition on Youth Voter Turnout." *Political Behavior* 30 (4): 415–36. DOI: 10.1007/s11109-008-9057-x. [9]

Palfrey, Thomas R., and Howard Rosenthal. 1983. "A Strategic Calculus of Voting." *Public Choice* 41 (1): 7–53. DOI: 10.1007/BF00124048. [10]

Plutzer, Eric. 2002. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96 (1): 41–56. DOI: 10.1017/S0003055402004227. [9, 30]

Pons, Vincent. 2018. "Will a Five-Minute Discussion Change Your Mind? A Countrywide Experiment on Voter Choice in France." *American Economic Review* 108 (6): 1322–63. DOI: 10.1257/aer.20160524. [8]

Pons, Vincent, and Guillaume Liegey. 2019. "Increasing the Electoral Participation of Immigrants: Experimental Evidence from France." *Economic Journal* 129 (617): 481–508. DOI: 10.1111/ecoj.12584. [8]

Pons, Vincent, and Clémence Tricaud. 2018. "Expressive Voting and Its Cost: Evidence from Runoffs with Two or Three Candidates." *Econometrica* 86 (5): 1621–49. DOI: 10.3982/ECTA15373. [8, 28]

Roth, Christopher, and Johannes Wohlfart. 2018. "Experienced Inequality and Preferences for Redistribution." *Journal of Public Economics* 167: 251–62. DOI: 10.1016/j.jpubeco.2018.09.012. [8, 10, 15, 16, 35]

Sears, David O., and Carolyn L. Funk. 1999. "Evidence of the Long-Term Persistence of Adults' Political Predispositions." *Journal of Politics* 61 (1): 1–28. DOI: 10.2307/2647773. [10]

Settle, Jaime E., Robert M. Bond, Lorenzo Coviello, Christopher J. Fariss, James H. Fowler, and Jason J. Jones. 2016. "From Posting to Voting: The Effects of Political Competition on Online Political Engagement." *Political Science Research and Methods* 4 (2): 361–78. DOI: 10.1017/psrm.2015.1. [11]

Severen, Christopher, and Arthur A. van Benthem. 2022. "Formative Experiences and the Price of Gasoline." *American Economic Journal: Applied Economics* 14 (2): 256–84. DOI: 10.1257/app.20200407. [8]

Shachar, Ron, and Barry Nalebuff. 1999. "Follow the Leader: Theory and Evidence on Political Participation." *American Economic Review* 89 (3): 525–47. DOI: 10.1257/aer.89.3.525. [10]

Shaw, Daron R. 1999. "The Methods behind the Madness: Presidential Electoral College Strategies, 1988–1996." *Journal of Politics* 61 (4): 893–913. DOI: 10.2307/2647547. [11]

Shue, Kelly, and Erzo F. P. Luttmer. 2009. "Who Misvotes? The Effect of Differential Cognition Costs on Election Outcomes." *American Economic Journal: Economic Policy* 1 (1): 229–57. DOI: 10.1257/pol.1.1.229. [8, 23]

Slotwinski, Michaela, and Alois Stutzer. 2018. "Women Leaving the Playpen: The Emancipating Role of Female Suffrage." CESifo Working Paper No. 7002. Munich: CESifo. URL: https://www.cesifo.org/DocDL/cesifo1_wp7002.pdf. [8]

Smets, Kaat. 2012. "A Widening Generational Divide? The Age Gap in Voter Turnout through Time and Space." *Journal of Elections, Public Opinion and Parties* 22 (4): 407–30. DOI: 10.1080/17457289.2012.728221. [9]

Smets, Kaat, and Anja Neundorf. 2014. "The Hierarchies of Age-Period-Cohort Research: Political Context and the Development of Generational Turnout Patterns." *Electoral Studies* 33: 41–51. DOI: 10.1016/j.electstud.2013.06.009. [9]

Spear, Linda Patia. 2000. "Neurobehavioral Changes in Adolescence." *Current Directions in Psychological Science* 9 (4): 111–14. DOI: 10.1111/1467-8721.00072. [10]

Spenkuch, Jörg L., and David Toniatti. 2018. "Political Advertising and Election Results." *Quarterly Journal of Economics* 133 (4): 1981–2036. DOI: 10.1093/qje/qjy010. [8]

Strömberg, David. 2008. "How the Electoral College Influences Campaigns and Policy: The Probability of Being Florida." *American Economic Review* 98 (3): 769–807. DOI: 10.1257/aer.98.3.769. [11]

Svaleryd, Helena, and Jonas Vlachos. 2009. "Political Rents in a Non-corrupt Democracy." *Journal of Public Economics* 93 (3–4): 355–72. DOI: 10.1016/j.jpubeco.2008.10.008. [7]

Taylor, Paul, Rich Morin, D'Vera Cohn, and Wendy Wang. 2008. "American Mobility: Who Moves? Who Stays Put? Where's Home?" A Social & Demographic Trends Report. Washington, DC: Pew Research Center. URL: https://www.pewresearch.org/wp-content/uploads/sites/3/2011/04/American-Mobility-Report-updated-12-29-08.pdf. [13]

The American National Election Studies. 1999a. "ANES 1952 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999b. "ANES 1956 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999c. "ANES 1958 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999d. "ANES 1960 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999e. "ANES 1964 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999f. "ANES 1966 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999g. "ANES 1968 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999h. "ANES 1970 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999i. "ANES 1972 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999j. "ANES 1974 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999k. "ANES 1976 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999l. "ANES 1978 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999m. "ANES 1980 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999n. "ANES 1982 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999o. "ANES 1984 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999p. "ANES 1986 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999q. "ANES 1988 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 22, 2019. [12]

The American National Election Studies. 1999r. "ANES 1990 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 23, 2019. [12]

The American National Election Studies. 1999s. "ANES 1992 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 23, 2019. [12]

The American National Election Studies. 2002. "ANES 1994 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 23, 2019. [12]

The American National Election Studies. 2003. "ANES 1996 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 23, 2019. [12]

The American National Election Studies. 2005a. "ANES 2000 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 23, 2019. [12, 36]

The American National Election Studies. 2005b. "ANES 2004 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed October 12, 2019. [12]

The American National Election Studies. 2015. "ANES 2008 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed December 17, 2019. [12]

The American National Election Studies. 2016. "ANES 2012 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed October 13, 2019. [12, 36]

The American National Election Studies. 2019a. "ANES 2016 Time Series Study [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed October 13, 2019. [12, 36]

The American National Election Studies. 2019b. "ANES Time Series Cumulative Data File (1948–2016) [dataset]." Stanford University and the University of Michigan [producers and distributors]. Available at https://www.electionstudies.org. Accessed September 27, 2019. [12]

U.S. Bureau of Economic Analysis. 2019. "Annual Personal Income and Employment by State: SAINC1 - Personal Income [online]." Available at https://apps.bea.gov/iTable/iTable.cfm?reqid=70&step=30&isuri=1&major_area=0&area=xx&year=-1&tableid=21&category=421&area_type=0&year_end=-1&classification=non-industry&state=0&statistic=3&yearbegin=-1&unit_of_measure=levels. Accessed December 3, 2019. [36]

U.S. Bureau of Labor Statistics. 2019. "CPI for All Urban Consumers: Series Id CUUR0000SA0 [online]." Available at https://data.bls.gov/timeseries/CUUR0000SA0. Accessed December 10, 2019. [36]

U.S. Census Bureau. 2017. "Historical Census of Housing Tables: Homeownership [online]." Available at https://www.census.gov/data/tables/time-series/dec/coh-owner.html. Accessed December 10, 2019. [36]

U.S. Federal Bureau of Investigation Uniform Crime Reporting Statistics. 2019. "UCR Data Online [online]." Prepared by the National Archive of Criminal Justice Data. Available at https://www.ucrdatatool.gov/. Accessed December 4, 2019. [37]

Verba, Sidney, Kay Lehman Schlozman, and Nancy Burns. 2005. "Family Ties: Understanding the Intergenerational Transmission of Participation." In *The Social Logic of Politics: Personal Networks as Contexts for Political Behavior*. Edited by Alan S. Zuckerman. Philadelphia, PA: Temple University Press. Chapter 5, 95–114. [27]

Vowles, Jack, Gabriel Katz, and Daniel Stevens. 2017. "Electoral Competitiveness and Turnout in British Elections, 1964–2010." *Political Science Research and Methods* 5 (4): 775–94. DOI: 10.1017/psrm.2015.67. [9]

Watanabe, Aaron. 2017. "US Census Bureau State Level Population Estimates, 1990–2016 [online]." Available at https://www.scholar.harvard.edu/awatanabe/data. Accessed December 8, 2019. [36]

Williams, Jhacova. Forthcoming. "Historical Lynchings and the Contemporary Voting Behavior of Blacks." *American Economic Journal: Applied Economics*, URL: https://www.aeaweb.org/articles?id=10.1257/app.20190549. [8]

Woolley, John T., and Gerhard Peters. 2019. "The American Presidency Project [online]." University of California, Santa Barbara [hosted]. Available at https://www.presidency.ucsb.edu/statistics/elections. Accessed December 15, 2019. [12]

Chapter 2

The Cost of Worrying about an Epidemic: Ebola Concern and Cognitive Function in the US*

2.1 Introduction

During an epidemic, the disease often does not come alone. It is accompanied by powerful emotional responses, manifesting in feelings of worry and fear. For instance, the share of adults in the US who reported experiencing significant worry on the previous day increased from 38 to 58 percent in March 2020 during the onset of the COVID-19 pandemic (Witters and Harters, 2020). Economists have argued that these emotions can affect behavior in important ways, e.g., by inducing temporary visceral urges to withdraw or negative anticipatory utility associated with the object of fear (e.g., Elster, 1998; Loewenstein, 2000; Caplin and Leahy, 2001). In line with these theoretical accounts, existing research on the economic consequences of epidemics emphasizes fear of contagion as a key channel that can affect patterns of economic activity by triggering avoidance behavior (e.g., Goolsbee and Syverson, 2021).

However, fear and worry might not only affect economic outcomes through deliberate changes of behavior. In this chapter, I empirically investigate another, hitherto unexplored potential consequence of worrying about epidemics: it might have a direct negative effect on cognitive function, an important determinant of economic decision-making and productivity (Heckman, Stixrud, and Urzua, 2006; Burks, Carpenter, Goette, and Rustichini, 2009; Dohmen, Falk, Huffman, and Sunde, 2010; Benjamin, Brown, and Shapiro, 2013). My hypothesis is based on studies from cognitive psychology, which suggest that anxiety can shift scarce attentional resources towards the perceived threat, thus reducing the available cognitive capacity for threat-unrelated processing tasks (Mathews, 1990; Eysenck, Derakshan, Santos, and Calvo, 2007; Robinson, Vytal, Cornwell, and Grillon, 2013; Moran, 2016; Sari, Koster, and Derakshan,

* I thank Lorenz Goette, Luca Henkel, and Florian Zimmermann for valuable comments and discussions. The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

2017). Worrying is associated with the experience of intrusive negative thoughts, which can induce cognitive load in the same way as disruptions due to external distractors like a honking car or an incoming phone call.[1] While an increase in focus on the threat may lead to a better response in situations with a high risk of infection, my aim in this chapter is to assess whether this comes at a cognitive cost in other domains. Since many of the decisions we take and tasks we perform during an epidemic are unrelated to the threat of the disease, the economic consequences could be substantial. In particular, a decline in available cognitive resources in threat-unrelated domains implies suboptimal economic choices and reduced labor productivity. Moreover, this cognitive cost of worrying might have a greater incidence than the disease itself because not everyone who worries also becomes infected and worry is often disproportionate to the actual threat.

I test this hypothesis in the context of an episode of heightened public concern about the possibility of a local Ebola outbreak in the US in October 2014. Ebola is a life-threatening hemorrhagic fever caused by Ebola virus, which emerges periodically in local outbreaks on the African continent. In the wake of the largest outbreak to date, four cases of Ebola were diagnosed in the US and led to the implementation of contact tracing procedures in the cities of Dallas, Texas, Cleveland, Ohio, and New York, New York throughout October 2014. The US Ebola cases were imported by incoming air travelers from affected African countries and spurred considerable public concern, with almost half of all respondents of a mid-October representative survey reporting worries that they or a family member would contract Ebola. Fears varied in intensity across space and time until subsiding in early November, as people learned that the US cases were well-contained by the authorities and did not lead to further local infections. One attractive feature of this setting is that the fraction of the US population that is physically affected by Ebola cases or resulting containment measures is vanishingly small, so any effects on cognition can only be driven by psychological responses to the disease outbreak.

For my analysis, I use data from tests of fluid intelligence, which were administered as part of a wave of survey interviews by the Health and Retirement Study (HRS). The HRS is a biennial panel study that conducts interviews with a representative sample of the US population aged 50 and older and covers a wide range of topics. Combining public and restricted HRS data sources, I obtain a sample of cognitive test scores from about 500 HRS interviews conducted in October 2014, with detailed information on their date and location. This allows matching test scores with a measure of the local level of Ebola concern test takers were exposed to in the week and media market of their interview, which I construct from data on search interest for the term "Ebola" from Google Trends.

In the first part of my analysis, I investigate the association between Ebola concern at the interview and the resulting cognitive test performance in ordinary least squares (OLS) regressions with various control variables. A major concern for observational studies on the effect of epidemic-induced worry on cognitive function is reverse causality: if the propensity

1. See, e.g., Deck and Jahedi (2015) for evidence that cognitive load impairs cognitive ability. Furthermore, other channels could also be at play. Notably, epidemics might increase emotional arousal more generally or evoke a scarcity mindset, with similar negative effects on cognitive function (e.g., Kaufman, 1999; Mullainathan and Shafir, 2013).

to worry about an epidemic is partly determined by education or cognitive ability, living in a media market with a high level of Ebola concern could be an indication of low cognitive function rather than a cause of it. To alleviate this concern, I control for demographic predictors of cognitive function like age and education, changes in life circumstances that could be related to cognitive decline like retirement, as well as cognitive test scores from previous HRS interview waves. Moreover, I add proxies for time-invariant characteristics of the interview location like the local search interest for the search topic "virus" in 2013, and I include interview week and Census region dummies. I find a strong and highly significant negative relationship between the level of Ebola concern HRS respondents are exposed to during their interview and their resulting cognitive test performance. In the most demanding specification, a one-standard-deviation (SD) increase in Ebola concern is associated with a decline in test scores of about 0.14 SD.

In the second part of the analysis, I use an instrumental variables (IV) strategy to strengthen the claim that the documented relationship is causal. This identification strategy exploits the fact that the specific timing and geographic location of the US cases has a random component. My approach builds on Campante, Depetris-Chauvín, and Durante (2020), who show that regional differences in Ebola concern are predicted by geographical distance to the closest US Ebola location and use this correlation for an IV estimation of the effect of fear on the outcomes of the 2014 US midterm elections. I demonstrate the existence of a similar first-stage relationship in my sample of HRS interviews, with predictably higher levels of Ebola concern in closer proximity to a publicly known US Ebola case. This holds also when controlling for distance to close large US cities more generally, thus accounting for the higher likelihood of Ebola case imports in urban centers. With respect to the exclusion restriction, I conduct two falsification tests. First, I show that in my sample of October interviews, the instrument is unrelated to the level of Ebola search interest in August—before the first US case—and to an equivalent search interest measure during the H1N1 pandemic in 2010. This alleviates the specific concern that distance to the closest US Ebola location could be associated with systematic regional differences in attitudes towards diseases or a tendency towards internet searches during disease outbreaks. Second, I find that there is no correlation between test scores and placebo versions of the instrument in the form of backdated US Ebola case dates in a different sample of HRS interviews conducted in September, just before the actual emergence of the first US Ebola case. These findings suggest that distance to the closest US Ebola location is a valid instrument for Ebola concern in my setting. The corresponding reduced-form and IV estimates support the conclusions of the OLS analysis and indicate a substantial cognitive cost of Ebola concern. Extrapolating from the point estimate of the IV coefficient, the increase in Ebola concern caused by the emergence of the US Ebola cases in October relative to the level of Ebola search interest in August, when the World Health Organization (WHO) declared the Ebola outbreak in Western Africa a public health emergency, would imply a reduction in test scores of about 12 scale points. By construction of the score scale, this corresponds to a drop in the probability of answering a test item of equivalent difficulty correctly from 90 to 71 percent. This effect survives various robustness checks, including different variants of the instrument and different approaches to constructing an Ebola concern measure from search interest data.

A remaining concern is that the composition of my sample of test scores may be distorted because potential HRS respondents react to nearby US Ebola cases for fear of contracting the disease. If individuals with positive unobservable shocks to cognitive ability in 2014 were more likely to delay their interview after the occurrence of a close case, the documented negative influence of Ebola concern on cognitive test performance could also be driven by a selection effect. To assess this possibility empirically, I test whether there is a differential effect of distance to the closest US Ebola location on the date of the interview for respondents with high and low predicted test scores, where predictions are based on all available observable determinants of cognitive ability in the HRS data. While I find that interviews in closer proximity to US Ebola locations are on average conducted a few days later, the estimates do not provide any indication that this relationship varies by cognitive ability. Therefore, the results of my analyses point to a direct causal effect of epidemic concern on cognitive function.

My findings relate to a large body of research on the economic consequences of epidemics. This topic has gained traction during the recent COVID-19 pandemic because understanding the specific mechanisms by which the economy is disrupted is crucial to effectively alleviate the adverse economic effects through targeted government measures.[2] Existing research highlights fear of contagion as a key channel that can affect patterns of economic activity by triggering avoidance behavior, thereby causing both supply and demand shocks (e.g., Eichenbaum, Rebelo, and Trabandt, 2021; Farboodi, Jarosch, and Shimer, 2021). For instance, consumers avoid spending on goods and services that involve interpersonal contact like travel, restaurant visits, or physical retail shopping (e.g., Rassy and Smith, 2013; Cox, Ganong, Noel, Vavra, Wong, et al., 2020; Chen, Qian, and Wen, 2021; Goolsbee and Syverson, 2021), and workers reduce their infection risk exposure by cutting labor supply or working from home (e.g., Brynjolfsson, Horton, Ozimek, Rock, Sharma, et al., 2020).[3] Moreover, there is evidence that drops in consumer expectations about future economic conditions during an epidemic are associated with fears of the disease (Binder, 2020; Fetzer, Hensel, Hermle, and Roth, 2021). These findings are in line with theoretical accounts on the influence of emotions on economic choices and judgments (Elster, 1998; Loewenstein, 2000), which emphasize the role of experienced emotions as motivators of specific mitigation behaviors, like an urge to withdraw from a dangerous situation in the case of fear. This perspective contends that emotional responses can cause a transient increase in the relative marginal benefit of specific choice alternatives relative to others, driven in part by emotion-induced changes in the weighting and perception of subjective probabilities, the evaluation of possible outcomes, and an elevated discounting of the future.[4] In contrast, I document an adverse effect of epidemic-induced worry on task-available cognitive

2. Here, I restrict attention to studies that concern the psychological aspects of epidemics. Brodeur, Gray, Islam, and Bhuiyan (2021) provide an overview of research on the economic effects of the COVID-19 pandemic with a wider focus. For other studies on the political, social, and economic effects of the 2014 Ebola epidemic, see Flückiger, Ludwig, and Önder (2019) and Campante, Depetris-Chauvín, and Durante (2020), González-Torres and Esposito (2020), and Kostova, Cassell, Redd, Williams, Singh, et al. (2019).

3. However, Balgova, Trenkle, Zimpelmann, and Pestel (2022) only find limited evidence of a relationship between job search intensity and health concerns during the COVID-19 pandemic in an analysis based on survey data from the Netherlands.

4. Also see Loewenstein, Weber, Hsee, and Welch (2001), who develop a psychological model of choice under risk that explains why emotional reactions to risky situations and subsequent behavior can be disproportionate to

capacity, which is a component of the decision-making process. Thus, my results point towards a general reduction in the quality of cognitive choice rather than a shift towards specific choice options. Since the cognitive function test underlying the main outcome variable of my analysis is completely unrelated to respondents' object of concern, changes in the relative desirability of emotion-mitigating choice options cannot explain the observed effect. Instead, my results are consistent with an analogous emotion-induced shift towards the object of concern on the level of cognitive or attentional resources. This opens up an alternative explanation for seemingly irrational observed behavior like the hoarding of goods during the COVID-19 pandemic (Baker, Farrokhnia, Meyer, Pagel, and Yannelis, 2020), which might be triggered in part by cognitive errors. It also suggests a specific new mechanism by which epidemics may disrupt the economy. Importantly, worry may affect the decision-making and productivity of all economic agents and in all domains, even if they are not physically affected by the disease itself or resulting government restrictions.

In independent research, Bogliacino, Codagnone, Montealegre, Folkvord, Gómez, et al. (2021) show that self-reported experience of various negative shocks during the COVID-19 pandemic is related to lower scores on a cognitive test. Because their analysis is correlational, they cannot rule out that this finding is driven by endogeneity in the sense that individuals with exogenously lower cognitive ability are more likely to be affected by negative shocks.[5] An additional priming intervention they conduct to substantiate a causal interpretation is ineffective. In contrast, my identification strategy relies on quasi-random variation in the date and location of Ebola cases in the US, and I provide evidence that selection effects are unlikely to play a role for my results.

My findings also contribute to a strand of literature in behavioral development economics on the psychology of poverty. In an influential book, Mullainathan and Shafir (2013) hypothesize that poverty itself can impair decision-making because perceived scarcity captures and taxes cognitive resources, resulting in a vicious cycle of poverty.[6] This idea finds support in results from priming experiments as well as analyses of natural variation in rainfall and income before and after payday for farmers in developing countries, which indicate that concerns about low levels of income and income uncertainty can reduce performance on cognitive tests (Mani, Mullainathan, Shafir, and Zhao, 2013; Lichand and Mani, 2020).[7] I extend the scope of this line of work by providing initial evidence that concerns about health can impede cognition in

the cognitive evaluation of the same risk. Their model and the reviewed evidence illustrate the various ways by which emotions can affect economic choices and judgments.

5. For instance, Adams-Prassl, Boneva, Golin, and Rauh (2020) document that less educated workers were more likely to lose their job during the early stages of the COVID-19 pandemic in the US and the UK. With respect to health shocks, Benson, Amato, Osler, Hosmer, and Malhotra (2021) find that the high school dropout rate is one of the strongest predictors of observed variation in COVID-19 cases and deaths across US counties.

6. See Dean, Schilbach, and Schofield (2019) and Kremer, Rao, and Schilbach (2019) for recent reviews that also discuss other dimensions of poverty that may impair cognition, like noise, environmental pollution, nutrition, sleep deprivation, and mental health. Haushofer and Fehr (2014) survey a related body of research that explores how poverty-induced increases in stress and negative affect impact preferences, which also implies changes in economic decision-making.

7. Recent investigations of the effect of financial strain on cognitive function yield mixed results that do not always replicate the findings of the initial studies (e.g., Carvalho, Meier, and Wang, 2016; Duquennois, 2022). The

similar ways as financial strain. This suggests that the higher incidence of diseases and their more negative health consequences in poor countries could also be a factor for documented puzzles in the economic behavior of the poor (e.g., Banerjee and Duflo, 2007).

The remainder of the chapter is structured as follows. In Section 2.2, I briefly outline relevant aspects of the Western African Ebola virus epidemic of 2014 and public reactions to it in the US, which provides the context of the study. Afterwards, I turn to the different data sources and methodological considerations for the construction of important analysis variables in Section 2.3. The results of an initial OLS analysis of the association between Ebola concern and cognitive function are presented in Section 2.4. In Section 2.5, I describe the identification strategy and findings of the IV analysis which forms the core of the chapter. I assess the possibility that my results are affected by selection into interview dates in Section 2.6. Finally, Section 2.7 discusses the interpretation of my main findings, and Section 2.8 concludes.

2.2 Background

The context of my analysis is an episode of heightened public concern about Ebola in the US in October 2014, after a few isolated cases emerged in the country as a consequence of the Western African Ebola virus epidemic. As outlined by Goeijenbier, van Kampen, Reusken, Koopmans, and van Gorp (2014), Ebola is a life-threatening hemorrhagic fever caused by Ebola virus. It is known for its high case fatality rate of between 50 and 90 percent and its frightening symptomatology. Typically, infected persons develop fever, vomiting, and diarrhea about four to ten days after exposure to the virus, followed by manifestations of internal and external bleeding. In lethal cases, death occurs due to circulatory shock, low blood pressure, or multi-organ failure about one to two weeks after the onset of initial symptoms. Even though chances of survival can be enhanced by early symptomatic treatment like the replacement of lost body fluids, no approved antiviral medication or vaccine for Ebola existed in the US in 2014. However, despite the severity of the disease for the infected, the objective risk of death from Ebola for individuals in developed countries is low. This is because human-to-human transmission requires direct contact with body fluids of symptomatic patients, making local outbreaks highly unlikely.

The Western African Ebola virus epidemic is the largest Ebola outbreak to date, causing a total of 11 310 reported deaths (World Health Organization, 2016). It began in a rural area in Guinea in December 2013 and spread rapidly across Guinea, Liberia, and Sierra Leone in the following year, despite multinational efforts aimed at containing transmissions. Coltart, Lindsey, Ghinai, Johnson, and Heymann (2017) identify traditional burial practices involving physical contact with the bodies of the deceased, limited healthcare capacities with inadequate protective equipment of healthcare workers, and a fast transmission of the virus into densely populated centers due to highly mobile communities as key factors that propagated the spread of the virus during this phase of the epidemic. On August 8, 2014, the WHO declared Ebola an international public health emergency (WHO Ebola Response Team, 2014). Still, the outbreak was contained to West Africa until the official end of the epidemic in 2016, with only a few

general state of the evidence is discussed in a recent review of the scarcity hypothesis by de Bruijn and Antonides (2022).

isolated cases imported into other regions of the world by returning travelers from the three affected countries.

Four Ebola cases were diagnosed on US soil, all in close succession during autumn 2014. As a result, the Centers for Disease Control and Prevention (CDC) implemented contact tracing procedures at three locations within the US to identify and monitor about 450 individuals at risk of exposure to the disease.[8] The first patient was an incoming traveler from Liberia who arrived in Dallas, Texas, on September 20 (Chevalier, Chung, Smith, Weil, Hughes, et al., 2014). He presented himself in the emergency department of a local hospital with fever symptoms twice in the following days but was only tested for Ebola on his second appearance, resulting in the first Ebola virus infection diagnosed in the US on September 30. Consequently, the patient was isolated for further treatment under special precautions, and potentially exposed contacts were traced and monitored by the CDC. The man died from Ebola on October 8. Two nurses involved in direct care of the first patient developed symptoms and were tested positive for Ebola in Dallas, Texas, on October 11 and 15, respectively. They constituted the second and third US Ebola case, leading to the tracing and monitoring of additional contact persons. Since the third patient had traveled to Cleveland, Ohio, in the days before her Ebola diagnosis while potentially being infectious, the CDC also implemented contact tracing procedures and conducted Ebola preparedness assessments of local hospitals at this location (McCarty, Basler, Karwowski, Erme, Nixon, et al., 2014). The fourth US case was a physician who returned to the City of New York, New York, after having treated Ebola patients in Guinea (Yacisin, Balter, Fine, Weiss, Ackelsberg, et al., 2015). He reported fever symptoms and tested positive for Ebola on October 23, resulting in the home confinement and monitoring of three close contact persons. The patient survived and was discharged from hospital on November 10.

Despite the occurrence of local Ebola cases, the objective risk of an epidemic outbreak in the US was considered to be extremely low by experts at that time (e.g., Gomes, Pastore y Piontti, Rossi, Chao, Longini, et al., 2014). Consistent with this assessment, there were no locally transmitted Ebola cases outside of the healthcare sector. Yet, the four existing US Ebola cases caused considerable public concern about a major Ebola outbreak in the US. In nationally representative surveys conducted in the second week of October, almost half of the respondents reported being worried that they or a family member would contract Ebola (SteelFisher, Blendon, and Lasala-Blanco, 2015). Fears were also reflected in disproportionate behavioral reactions like private hazmat suit purchases and the shunning of returning travelers from Africa even when they had visited countries without documented Ebola cases (Bedrosian et al., 2016). SteelFisher, Blendon, and Lasala-Blanco (2015) conjecture that public concern was fueled by sensationalist media coverage, limited trust in the US federal government and health authorities, and misperceptions about the contagiosity of the disease, with a substantial fraction of the

8. Six additional US citizens were diagnosed with Ebola in 2014 while working with medical teams to stop the epidemic in West Africa. These patients were medically evacuated to the US for treatment in one of four specialized hospitals with biocontainment units (Rainisch, Asher, George, Clay, Smith, et al., 2015) and did not pose an infection risk for the US population at any time. Although the evacuations also received media attention, they did not provoke a comparably strong emotional response. For instance, they resulted in significantly fewer e-mail inquiries and visits to the CDC's Ebola webpages than the US-diagnosed cases (Bedrosian, Young, Smith, Cox, Manning, et al., 2016).

population in the belief that Ebola is airborne.[9] Eventually, public concern subsided as people learned that isolated Ebola cases in the US were well-contained and did not lead to further local infections. Consequently, the level of worry about Ebola reported in national surveys saw a marked decline in early November (SteelFisher, Blendon, and Lasala-Blanco, 2015).

The episode of high public concern during October 2014 poses an ideal setting for my analysis of the effects of emotional responses to epidemic risk on cognitive function. First, as illustrated by Campante, Depetris-Chauvín, and Durante (2020), it is characterized by significant geographic and temporal variation in worry about Ebola, with predictably higher concern in close proximity to recent US Ebola cases. Second, the specific timing and geographic location of US cases has a random component. In contrast to outbreaks that originate within a country because of specific characteristics of the local environment or population, the US Ebola cases were imported by incoming air travelers. Due to minor coincidences, infected travelers could just as well have arrived on a different day in another large US city, resulting in a very different case pattern. This makes it unlikely that proximity to US Ebola cases is systematically related to regional differences in cognitive ability after accounting for distance to large US cities, thus enabling an IV strategy to identify a causal effect. Third, public concern was disproportionate to the actual level of epidemic activity. The fraction of the US population that was physically affected by the disease or resulting containment measures is vanishingly small. There were also no financial ramifications for the average citizen. This implies that US Ebola cases can only affect cognitive function via the public's emotional response.

2.3 Data

To test the hypothesis that epidemic-induced worry impedes cognitive function in the context of the occurrence of Ebola cases in the US, I require a dataset with geographically disaggregated measures of cognitive function and Ebola concern in the US during October 2014. Since temporal and geographic variation in Ebola concern can be large, it is especially important that the data includes precise information on both the location and date of cognitive performance observations. To satisfy this requirement, my approach is to supplement a dataset of cognitive function test scores with a measure of Ebola concern constructed from Google Trends data.

2.3.1 Data on Cognitive Function

My main data source is the HRS, a panel study that conducts interviews about various topics related to health, economic situation, and family status with a representative sample of the US population aged 50 and older every two years. Importantly, it also includes tests of cognitive function. Interviews are carried out either by telephone or in person at respondents' homes, with the interview mode and date determined in part by respondent availability. For the 2014 wave of the HRS, 18 747 respondents were interviewed between March 2014 and April 2015.

9. Towers, Afzal, Bernal, Bliss, Brown, et al. (2015) provide evidence that Ebola-related news videos contributed to the propagation of fear based on parameter estimates from a model of news contagion.

Of these interviews, 801 were conducted in October, during the period of heightened Ebola concern described in Section 2.2. These October interviews are the main focus of my analysis.

I use data on cognition and various sociodemographic characteristics from the 2010 to 2014 waves taken from the RAND HRS Longitudinal File (RAND Center for the Study of Aging, 2020) as well as additional variables from the individual wave datasets RAND HRS 2010 Fat File (RAND Center for the Study of Aging, 2021a), RAND HRS 2012 Fat File (RAND Center for the Study of Aging, 2021b) and RAND HRS 2014 Fat File (RAND Center for the Study of Aging, 2021c). To pinpoint the circumstances of each interview, I additionally use restricted data on interview locations (Health and Retirement Study, 2019a) and interview dates (Health and Retirement Study, 2019b). In particular, I have access to information on the county in which the interview took place (item STCTYFIPS10) and the beginning date of the interview (items OIWMONTH, OIWDAY and OIWYEAR). A minor limitation of this data is that the beginning date of the interview does not always equal the date of the cognitive tests. A small fraction of interviews is suspended at some point and completed on a later date, e.g., because the respondent suffers from an acute health problem. I restrict my main analysis sample to interviews that were both started and completed in October 2014, but I do not observe whether or at what point an interview was interrupted within the month. Therefore, the date assigned to each interview is a lower bound for the day on which the cognitive tests were conducted. As a result, my analysis may falsely treat some cognitive function tests as having occurred before rather than after a given US Ebola case, which would typically imply an overestimation of the distance to the closest Ebola US case and an underestimation of the level of Ebola concern that test was subject to.

The HRS contains tests of three different aspects of fluid intelligence, the concept that economists often study when they are interested in cognitive ability: quantitative reasoning, verbal reasoning, and retrieval fluency (Fisher, McArdle, McCammon, Sonnega, and Weir, 2014).[10] A subset of the tests is administered in each wave. The only fluid intelligence test available in the 2014 wave for a large number of respondents is the verbal analogies task, a measure of verbal reasoning based on version III of the Woodcock-Johnson Tests of Cognitive Abilities. Each test item consists of a word pair that defines a specific logical relationship and a single third word. The respondent's task is to complete the analogy by naming the matching fourth word such that the logical relationship of the newly created pair is identical to that defined by the given word pair. Items are read out to the respondent in the form of open-ended sentences like "Mother is to Daughter as Father is to ..." without response options. In the HRS, the verbal analogies task is conducted as a block-adaptive test consisting of a total of six items. The difficulty of later items depends on respondents' performance on early items, allowing to get a relatively nuanced assessment of performance even with a short test.

The test result is reported as a W score of verbal reasoning, which constitutes my main outcome measure (item OVESCORE in the RAND HRS 2014 Fat File). The W score is defined on a numeric scale in terms of the difficulty of the item that a respondent is predicted to answer

10. For instance, Raven's Matrices is a test of fluid intelligence that has been used in the economic literature on the psychology of poverty (e.g., Mani et al., 2013). Fluid intelligence is defined as an individual's general ability to reason and solve problems in novel situations that do not depend on acquired knowledge.

correctly with exactly 50 percent probability (Jaffe, 2009). The scale is centered on a value of 500, which marks the performance of an average ten-year old. Importantly, the W score is designed as an interval score, implying that a given score difference always corresponds to the same performance difference, irrespective of where on the scale it is. For instance, a ten-point increase in W score always implies an increase in the probability of correctly solving the old score's reference item from 50 percent to 75 percent. This interval property is crucial for my purpose because it is an implicit assumption underlying the analysis of average treatment effects (Jacob and Rothstein, 2016).

The panel structure of the HRS also allows me to control for cognitive function test scores from the two previous survey waves, for which the set of eligible respondents was largely identical.[11] In these waves, the two other measures of fluid intelligence were administered: respondents were asked to solve number series tasks to test their quantitative reasoning, and their retrieval fluency was measured by the number of distinct animals they could name within a time limit of one minute. In addition, the HRS also contains cognitive tests designed to detect early signs of cognitive decline that often precede neurodegenerative diseases like dementia or Alzheimer's disease. Of these, I use the scores of a word recall test and a serial sevens test as additional controls.[12] Finally, the rich set of variables collected in the HRS also includes information on a number of demographic characteristics and on changes in general health status or lifestyle since the previous wave that could be associated with cognitive decline.

2.3.2 Measuring Ebola Concern

I construct measures of Ebola concern for the location and date of the HRS interviews based on data on internet search interest provided by Google Trends. Such data is nowadays frequently used in economic research to study behavior in real-time or to proxy for outcomes for which no direct measure is available (e.g. Da, Engelberg, and Gao, 2011; Baker and Fradkin, 2017; Bacher-Hicks, Goodman, and Mulhern, 2021; Brodeur, Clark, Fleche, and Powdthavee, 2021).

Google Trends provides data on the number of searches for a requested keyword or topic relative to the total Google search volume in a given region and time period. The data is normalized such that the highest relative search interest in the query has a value of 100. Two features of the data pose methodological challenges: First, no data is reported for region-periods for which the absolute number of searches of the keyword is below a specific threshold for reasons of privacy protection. Second, Google Trends only evaluates a random subsample of the universe of all searches, so the provided values can fluctuate, especially when data of high temporal and geographic granularity is requested for sparsely populated regions.

I obtain data on search interest for the keyword "Ebola" by week from August to October 2014 for all media markets in the US from the Google Trends website.[13] This is the highest available level of granularity with good data quality for this keyword. To smooth out the sam-

11. Refer to Appendix 2.A for details on the construction of control variables based on HRS data.

12. The other cognitive tests are either too easy, like naming the current date, or only collected for a subset of respondents in different waves.

13. Media markets are geographical areas—usually groups of counties—that receive the same television and radio signals. There are 210 media markets in the US.

pling error for smaller regions, I repeat the query 100 times for each media market. In each request, I also include the US as a whole and rescale the returned time series such that the peak of the US time series—the week from October 12 to 18—has a value of 100 to make the data from different queries comparable. Missing values resulting from Google's privacy threshold introduce a selection bias into these time series because data for some weeks in smaller media markets is only reported if a random sample features exceptionally high relative Ebola search interest for these media market–weeks. To deal with this source of bias, I code missing values as zero and then use the median across all iterations, including those with missing values. This results in an unbiased estimate of the median as long as data is reported in at least 51 iterations, which is the case for all relevant weeks and media markets for my sample of HRS respondents.[14] The resulting dataset is a media market–week panel of Ebola search interest, which I merge to the interview counties and dates of my HRS respondent sample using a crosswalk provided by Sood (2016).

In addition, I collect search interest data for the keyword "anxiety" and the topic "virus (infectious agent)" in all media markets for the year 2013, and for the topic "swine flu (swine influenza)" for the period between April 25, 2009—the day the WHO declared a public health emergency after the novel H1N1 influenza virus started to spread across several US states—and August 10, 2010—the day the WHO declared the H1N1 pandemic over.[15] Again, I repeat each query 100 times and use the median value for each media market. These variables are used as covariates or placebo outcomes in the analysis.

For my analysis, I want to construct a measure of the level of Ebola concern HRS respondents feel on the day of their interview, which is not necessarily equal to the local level of Ebola search interest during that week. Campante, Depetris-Chauvín, and Durante (2020) show in daily data that local Ebola search interest and Twitter activity react strongly to new US Ebola cases, but these responses are short-lived. In particular, they fade much faster than would be expected on the basis of the national surveys of Ebola worries reported in SteelFisher, Blendon, and Lasala-Blanco (2015). This suggests that Ebola concern lingers past initial search activity: individuals seem to actively search information about Ebola on the internet only when they first get worried about a nearby US Ebola case and afterwards rely on passive information acquisition through media reports, but they remain concerned. In my main specifications, I therefore use the average search interest in a given media market across all weeks between the week of the first US Ebola case and the interview week as my measure of Ebola concern. This closely follows the approach taken by Campante, Depetris-Chauvín, and Durante (2020), who use aggregate search interest across a five-week period between the first US case and the US midterm elections on November 4.

14. The results of a robustness check that uses another approach to data aggregation are presented in Section 2.5.4.

15. Topics are collections of Google searches that capture a specific meaning rather than an exact search phrase. I use topics rather than keywords for "virus" to avoid capturing the alternative meaning "computer virus" and for "swine flu" to also capture searches for the common alternative name "H1N1". For "anxiety", the distinction between keyword and topic is empirically irrelevant.

2.3.3 Other Data Sources

To compute the kilometer distance of interview counties to the closest publicly known US Ebola case on the date of the interview as well as to the closest large urban center, I rely on data from the US Census Bureau. I determine the 100 largest US cities by 2014 population using population estimates reported in U.S. Census Bureau (2020). Geographic coordinates of US Ebola locations and large US cities are taken from Gazetteer Files (U.S. Census Bureau, 2014), and I use coordinates of the mean centers of population of interview counties from U.S. Census Bureau (2019). All calculated distances are geodetic distances, i.e., they correspond to the length of the shortest curve along the surface of an ellipsoidal model of the earth, based on equations derived in Vincenty (1975). They are included in regression models in log-transformed form because changes in the distance to a US Ebola case should affect the level of Ebola concern in a given county less the farther away it is.

2.3.4 Main Sample

In addition to respondents with missing data on key variables, I also exclude respondents from Alaska and Hawaii because variation in geographical distance to Ebola cases in the contiguous US is unlikely to matter for their level of Ebola concern. This results in a final sample of 495 respondents from 108 media markets, covering all Census regions of the US. Descriptive statistics for this sample are reported in Table 2.B.1.

2.4 OLS Estimation

I start off the analysis by investigating the association between the level of Ebola concern HRS respondents are exposed to during their interview and their resulting cognitive test performance in OLS regressions with various control variables. This exploits all variation in Ebola concern across media markets and weeks observed in the wake of the occurrence of Ebola cases in the US in October 2014. The estimation equation is

$$\begin{aligned} \mathit{VerbalReasoningScore}_{ict} = {} & \alpha + \beta \mathit{EbolaConcern}_{mt} + \boldsymbol{\gamma}_1' \boldsymbol{X}_i + \boldsymbol{\gamma}_2' \boldsymbol{C}_c + \boldsymbol{\gamma}_3' \boldsymbol{M}_m + \delta_t + \rho_r \\ & + \epsilon_{imt}, \end{aligned} \tag{2.1}$$

where $\mathit{VerbalReasoningScore}_{ict}$ is the cognitive function test score individual i achieves in her HRS interview in county c and week t, and $\mathit{EbolaConcern}_{mt}$ is the level of Ebola concern in the corresponding media market in that week. The hypothesis that worrying about Ebola is associated with lower cognitive function corresponds to $\beta < 0$. Depending on the specification, I additionally include a vector of covariates $\boldsymbol{X}_i$, which can contain demographic variables, cognitive function test scores from previous HRS waves, and interview characteristics. Similarly, $\boldsymbol{C}_c$ and $\boldsymbol{M}_m$ are vectors of covariates on the county-level and media market–level, and δ_t and ρ_r are dummies for interview week and Census region, respectively.

All regressions use cluster-robust standard errors that account for heteroscedasticity and arbitrary intra-cluster correlation within media markets. This is necessary because the current level of Ebola concern in a given media market also depends on US Ebola cases from previous weeks, implying that Ebola concern is correlated across time within media markets.

The coefficient estimates of the OLS regression model are reported in the first four columns of Table 2.1. In column (1), I estimate Equation 2.1 without any controls, finding that being exposed to higher levels of Ebola concern during the interview is associated with significantly lower test scores. Of course, the documented relationship is susceptible to endogeneity problems. Especially in light of the potential role of misperceptions about the transmission dynamics of Ebola for the spread of fear, a major concern is reverse causality: media markets in which a higher fraction of the population is relatively less educated or has lower cognitive function might develop more worry about an Ebola epidemic. Then, living in a media market with a high level of Ebola concern would be an indication of low cognitive function rather than a cause of it. To mitigate this concern, I exploit the richness of the HRS data and add a variety of control variables.

Table 2.1. Ebola Concern and Cognitive Function: OLS Estimates

	Dependent variable: Verbal reasoning score				
	(1)	(2)	(3)	(4)	(5)
Ebola concern	−0.407***	−0.244***	−0.263***	−0.286***	
	(0.115)	(0.078)	(0.090)	(0.095)	
Ebola searches before first US case					0.088
					(0.294)
Demographic controls	No	Yes	Yes	Yes	No
Baseline cognition controls	No	Yes	Yes	Yes	No
Interview controls	No	No	Yes	Yes	No
Location controls	No	No	Yes	Yes	No
Census region dummies	No	No	No	Yes	No
Observations	495	495	495	495	495
R^2 (adjusted)	0.037	0.444	0.440	0.441	−0.002
Mean (dependent variable)	505.115	505.115	505.115	505.115	505.115

Notes: OLS estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern* is the average weekly relative search interest for the term “Ebola” in the media market of the interview location over all weeks between September 28 and the week of the interview. *Ebola searches before first US case* is the relative search interest for the term “Ebola” in the media market of the interview location in the first week of August. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics “anxiety” and “virus” in 2013. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

In column (2), I control for demographic characteristics that predict performance on cognitive tests, like age and education, as well as changes in life circumstances since the previous interview wave that could be associated with cognitive decline, like being diagnosed with dementia, moving into a nursing home, starting to live alone, or retiring from work. Moreover, I include several measures of baseline cognitive ability, using scores from tests of cognitive function in the previous two HRS interview waves. In particular, I add 2012 scores and changes in scores from 2010 to 2012 of the other two aspects of fluid intelligence tested in the HRS, quantitative reasoning and retrieval fluency. Analogously, I also include levels and trends of scores from a word recall and serial sevens test, both of which are aimed at detecting early warning signs of cognitive decline.

In column (3), I additionally include control variables for interview characteristics. These consist of the number of call attempts by the interviewer until an interview is conducted, a dummy for telephone interviews, and a set of dummies for the calendar week of the interview. This addresses a concern that arises as a result of the HRS field operations. Since interview mode and date are partly determined by participant responsiveness, those who are in worse health or go through personal turmoil during the data collection phase tend to be interviewed later during the wave. As a result, later interview dates can be a manifestation of unexpected negative health or cognitive capacity shocks that could introduce a spurious correlation between Ebola concern and cognitive test scores because both are correlated with interview dates.[16] With interview week dummies, the coefficient on Ebola concern is identified from geographic variation only and therefore immune to this concern. A further group of control variables is directed at time-invariant influences of the interview location. To capture systematic differences across media markets in general susceptibility to fear and the tendency to search for diseases on the internet, I add measures of search interest for "anxiety" and "virus" in 2013 as well as Ebola search interest before the first US case. As a result, the coefficient of interest in this specification isolates the effect of additional Ebola searches that arise from the perceived threat of an epidemic after the emergence of cases in the US. I also include the distance to the closest large US city because this is an important covariate in the IV estimation.

Finally, the fourth column adds Census region dummies, ruling out that an observed relationship is driven by arbitrary other regional determinants of test performance that might be correlated with Ebola concern.

The coefficient estimate shrinks after the inclusion of demographic and baseline cognition controls, but remains highly significant in all specifications. In the most demanding specification with the full set of covariates, the estimated reduction in verbal reasoning score is 4.1 W scale points, corresponding to about 0.138 SD, for a one SD increase in Ebola concern.

Comparing this effect size to that of other impediments of cognition is difficult because other studies employ different tests to assess different dimensions of cognitive functioning. Park (2022) reports a 0.055 SD decrease in student performance on standardized exams per SD increase in exam day temperature. In a field experiment, Dean (2021) finds that a ten decibel increase in engine noise reduces a summary index of performance on several cognitive function tests by 0.068 SD. Taken at face value, my coefficient estimate indicates that the effect of worrying about an epidemic is about twice as large as that of these environmental factors. On the other hand, it is markedly smaller than the 0.6 to 0.7 SD increase in Raven's Matrices

16. Indeed, a regression of cognitive function on interview date across all HRS respondents in the 2014 wave yields a significantly negative relationship.

scores of financially constrained sugarcane farmers after the receipt of their annual harvest income that can be derived from Mani et al. (2013).

For a different way of putting the magnitude of the effect into perspective, consider the implied difference in cognitive function between a situation with and without Ebola worries in the US. From Table 2.B.1, the difference between the average level of Ebola concern in my sample during October 2014 and average Ebola search interest in the first week of August, which constitutes the peak of the pre-October time series, is $63.8 - 27.3 = 36.5$ units.[17] Assuming that search interest in August is driven by public interest and the October increase stems from worrying about the possibility of a US epidemic, public concern about Ebola would be associated with a reduction in verbal reasoning score of 10.4 points on the W scale. This corresponds to a drop in the probability of answering a test item of equivalent difficulty correctly from 50 to about 25 percent (Jaffe, 2009).[18] This extrapolation indicates a substantial cognitive cost of worrying about an epidemic, with the potential of adverse economic consequences both in the form of suboptimal decisions in the private sphere and reduced productivity at work.

Column (5) of Table 2.1 contains estimates from a placebo check. In particular, I verify that there is no correlation between the cognitive function of my HRS respondent sample and Ebola search interest before the first US Ebola case, when geographic variation in search interest is unlikely to be driven by worry about an epidemic in the US. To this end, I regress the cognition test scores of October HRS interviews on Ebola search interest in the media market of the interview in the first week of August, when the WHO declared Ebola an international health emergency. This addresses the possibility that respondents with lower cognitive function happen to come from media markets with systematically higher public interest in Ebola or a higher propensity to search for Ebola on Google, for reasons other than worry. Such a correlation could confound the estimate of β in the regression model of Equation 2.1. Therefore, it is reassuring that the coefficient estimate of this exercise is close to zero and far from statistically significant.

2.5 IV Estimation

The previous section documents a robust and economically meaningful association between the level of Ebola concern individuals are exposed to at the time and location of a cognitive test and their resulting test scores. It also alleviates the most obvious endogeneity concerns. Yet, it is difficult to justify a causal interpretation of the observed relationship based on selection on observables alone. While controlling for individuals' test scores in 2010 and 2012 rules out reverse causality based on stable differences in cognitive ability, Ebola concern could still be related to unobserved local factors that accelerate cognitive decline. For instance, the strength of age-associated cognitive impairment and Ebola concern could both be correlated with local residents' media consumption, which would imply a downward bias. On the other

17. The equivalent numbers for the national US are very similar.

18. The corresponding drop in success probability is 15 percentage points for an easier test item that is answered correctly at a rate of 90 percent in the absence of concern about Ebola.

hand, measurement error in the independent variable might induce attenuation bias towards zero.

To overcome these remaining challenges, I implement an IV strategy that follows Campante, Depetris-Chauvín, and Durante (2020). The underlying idea is that people are more worried about Ebola if they live in close proximity to a recent US Ebola case. Since the specific location and timing of Ebola cases in large US cities is quasi-random, the distance to the actual cases then provides a source of exogenous variation in Ebola concern.

To exploit this, I construct a variable $EbolaDistance_{ct}$, which is the logarithm of the kilometer distance between the county of the HRS interview of individual i and the closest location with a relationship to a US Ebola case that was publicly known on interview day. The three Ebola locations are Dallas, Texas, where the first and second US Ebola case were diagnosed on September 30 and October 11, respectively, Cleveland, Ohio, which the third US Ebola case had visited immediately before her diagnosis on October 15, and the City of New York, New York, where the fourth and last US case was diagnosed on October 23. If (i) the relationship between the newly constructed variable and my measure of Ebola concern is strong and (ii) it satisfies the exclusion restriction, i.e., it is uncorrelated with other determinants of cognitive function after adjusting for covariates, then distance to the closest US Ebola location is a valid instrument for Ebola concern. Then, two-stage least squares estimation of Equation 2.1 yields an unbiased estimate of the causal effect of Ebola concern on cognitive function. In contrast to the OLS analysis, this identification strategy uses only the share of observed variation in Ebola concern that is the result of the exogenous shock posed by the emergence of close US Ebola cases.

To increase statistical precision and adjust for potential sources of correlation between the instrument and the outcome, I include the same control variables as in the OLS estimation in Section 2.4. Importantly, the covariates include distance to the closest large US city to account for a higher probability of a nearby Ebola case close to places with high travel volume. This control variable is constructed in the same way as the instrument, but considers the 100 largest US cities by 2014 population irrespective of their relationship to a US Ebola case.[19]

In the following subsections, I first show that distance to US Ebola cases is a strong predictor of Ebola concern. I then present two falsification tests conducted to assess the validity of the exclusion restriction. Afterwards, I report the results of the IV estimation and wrap up by summarizing the outcomes of various robustness checks.

2.5.1 Ebola Concern and Distance to US Ebola Locations

To convey a first impression of the relationship between proximity to Ebola cases and Ebola concern, the maps in Figure 2.1 depict the geographic distribution of Ebola concern across all media markets of the contiguous US. For each of the five weeks of October 2014 separately, deciles of the distribution of Ebola concern are colored in different shades of blue, with darker tones indicating media markets with higher levels of concern. The three Ebola locations appear

19. Note that the three US Ebola locations are themselves among the 100 largest cities in the US, taking the 1st, 9th, and 48th place in the ranking.

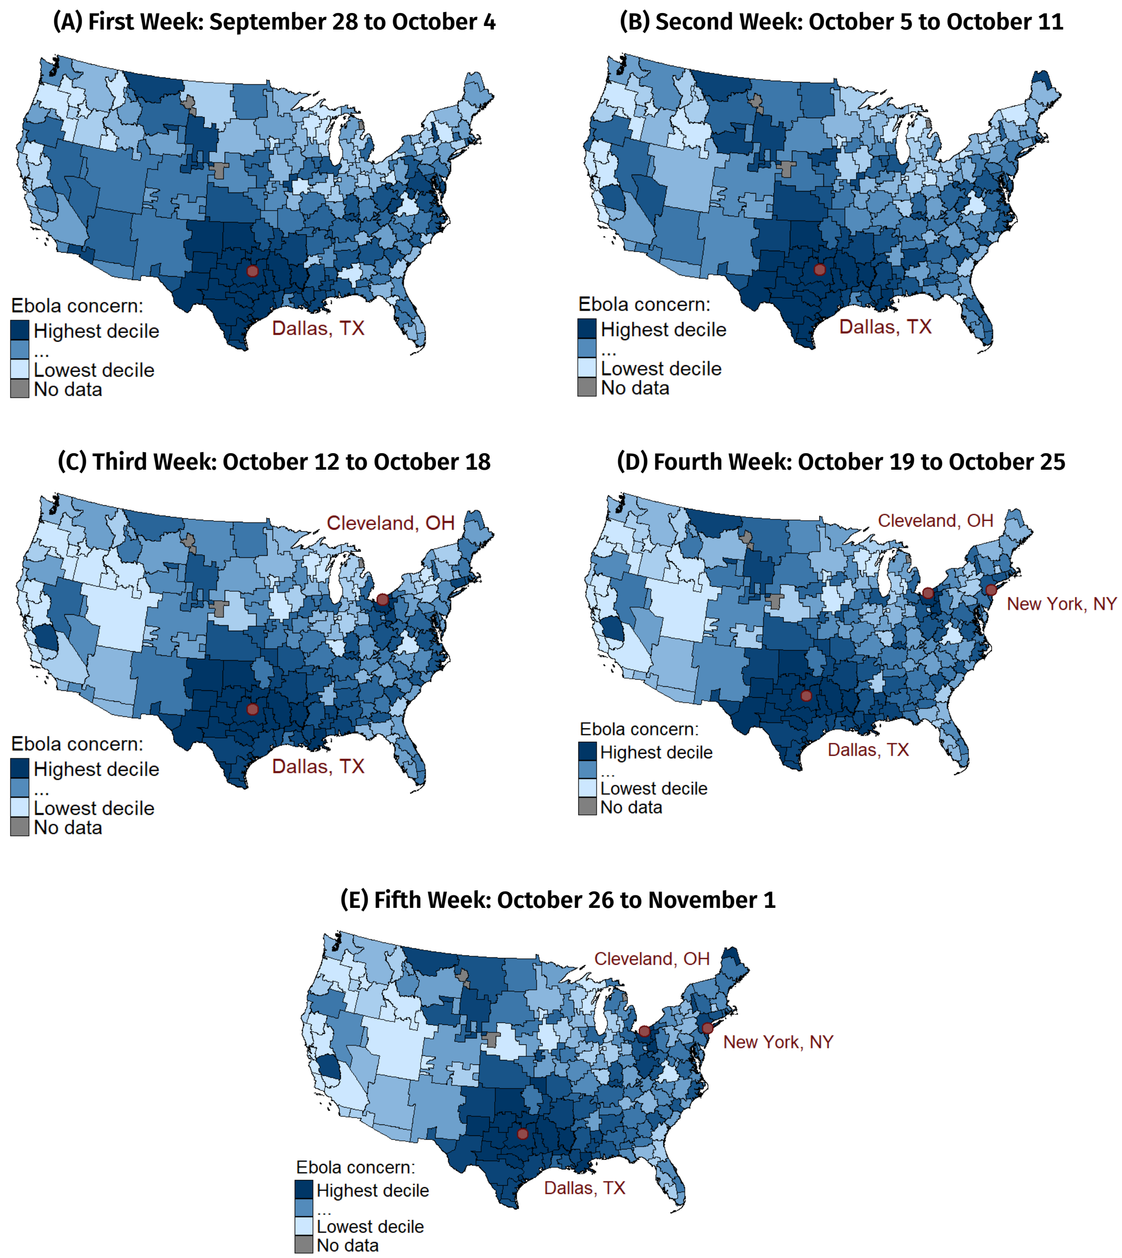

Figure 2.1. Ebola Case Locations and Geographic Variation in Ebola Concern across Media Markets by Week

Notes: Choropleth maps of the geographic variation in the main measure of Ebola concern for each week in October 2014. The level of Ebola concern in a given media market and week is defined as the average weekly relative search interest for the term "Ebola" in that media market over all weeks since September 28. Each panel shows the spatial distribution of Ebola concern within a specified week, with media markets in higher deciles colored in darker shades of blue. Media markets for which the Ebola concern measure cannot be constructed because of insufficient search interest data are colored in gray. Locations with a relationship to a US Ebola case that was publicly known by the end of that week are marked with red dots: The first and second US case were diagnosed in Dallas, TX, on September 30 and October 11, respectively. The third Ebola diagnosis, with connections to both Dallas, TX, and Cleveland, OH, became publicly known on October 15. The fourth and last US case was diagnosed in the City of New York, NY, on October 23.

on the maps as red dots starting from the week in which the respective case was diagnosed. Clearly, proximity to Ebola locations is associated with higher levels of worry in all weeks of October.

Table 2.2 contains estimates from corresponding first-stage regressions in my sample of HRS interview respondents, allowing for a formal assessment of the strength of the relationship between distance to close US Ebola locations and Ebola concern. As a natural benchmark, I first regress the level of Ebola concern in the media market of the interview on distance to the closest large US city. The estimated coefficient is displayed in column (1) of the table. It is close to zero and statistically insignificant, implying that distance to large cities alone does not predict Ebola concern.

Table 2.2. Ebola Concern and Distance to US Ebola Locations: First-Stage Estimates

	Dependent variable: Ebola concern			
	(1)	(2)	(3)	(4)
Distance to closest large city	0.161		0.555	−0.303
	(0.847)		(0.708)	(0.445)
Distance to closest Ebola location		−8.823***	−8.896***	−9.221***
		(2.756)	(2.741)	(2.602)
Demographic controls	No	No	No	Yes
Baseline cognition controls	No	No	No	Yes
Interview controls	No	No	No	Yes
Location controls	No	No	No	Yes
Census region dummies	No	No	No	Yes
Observations	495	495	495	495
R^2 (adjusted)	−0.002	0.303	0.305	0.659
Mean (dependent variable)	63.786	63.786	63.786	63.786
Effective F-statistic		10.249	10.532	12.563

Notes: First-stage estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern* is the average weekly relative search interest for the term “Ebola” in the media market of the interview location over all weeks between September 28 and the week of the interview. *Distance to closest large city* is the logarithm of the kilometer distance to the closest of the largest 100 US cities by 2014 population. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include Ebola concern before the first US Ebola case and media market–level relative search interest for the topics “anxiety” and “virus” in 2013. The effective F-statistic reported in the bottom of the table accounts for the use of cluster-robust standard errors (Montiel Olea and Pflueger, 2013). * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Columns (2) to (4) then present the first-stage estimates from regressions on the proposed instrument. I include distance to the closest US Ebola case once as the only independent variable, once only with adjustment for distance to the closest large city, and once with the full set of covariates. In all cases, the coefficient is large and highly significant. A one-unit increase in logarithmic distance to the closest US Ebola case, which is equivalent to 1.11 SD, increases Ebola concern by about nine units. This corresponds to one-fourth of the total increase in public concern about Ebola in October relative to August 2014. As indicated by the R^2 in column (2), proximity to US Ebola cases explains 30 percent of the variation in Ebola concern in my sample.

In the bottom row of the table, I report the effective F-statistic, which is a diagnostic tool to detect potential problems related to weak instruments in settings with cluster-robust standard errors (Montiel Olea and Pflueger, 2013). Weak instruments denote an insufficiently strong relationship between the instruments and the endogenous variable of interest. This is problematic because (i) it can induce an asymptotic bias that distorts the IV estimate towards the OLS estimate, and (ii) it can lead to size distortion in hypothesis tests on the IV coefficients, implying that the probability of falsely rejecting the null hypothesis is larger than the nominal size of the test. The effective F-statistic speaks to the likelihood of both of these problems.

For the preferred first-stage specification in column (4), the effective F-statistic is 12.563, which exceeds the standard Staiger and Stock (1997) rule-of-thumb threshold of ten. In simulations based on a sample of 230 recent IV specifications from the *American Economic Review*, Andrews, Stock, and Sun (2019) do not observe any bias or size distortions for specifications with effective F-statistics above this threshold, suggesting that my first stage is strong enough to rely on conventional hypothesis tests for inference. However, it is below the critical value suggested by Montiel Olea and Pflueger (2013) to formally reject an approximate asymptotic bias of ten percent relative to a "worst case benchmark" at the five percent level, which is at 23.109.[20] For my instrument, I can only reject a worst case bias of 30 percent, for which the critical value is 12.039.

With respect to test size distortions due to weak instruments, the effective F-statistic can be used along with the critical values of Stock and Yogo (2005) in the single-instrument case (Andrews, Stock, and Sun, 2019). Here, I can reject at the five percent level that the maximum size of a hypothesis test on the IV coefficient is more than 15 percent based on a Stock-Yogo critical value of 8.96, implying that the worst case size distortion is no more than ten percentage points. At the same time, I cannot reject that the maximum size of a hypothesis test on the coefficient is more than ten percent, for which the respective critical value is 16.38.

To sum up, while the strength of my instrument seems sufficient based on rules of thumb typically used in applied economics, the inability to reject intermediate levels of bias from a worst-case perspective leaves some room for concern. Therefore, I additionally report the results of IV inference procedures that are robust to weak instruments when presenting the IV estimation results.

20. Intuitively, the worst case benchmark used in Montiel Olea and Pflueger (2013) corresponds to a situation in which instruments are completely uninformative and first-stage and second-stage errors are perfectly correlated. The resulting critical values concern the null hypothesis that the approximate asymptotic bias of the IV estimator exceeds a fraction of the bias of this benchmark for at least one value of the parameter space.

2.5.2 Falsification Tests for the Exclusion Restriction

The validity of distance to the closest US Ebola case as an instrument for Ebola concern in Equation 2.1 relies on the exclusion restriction. It requires that the instrument is uncorrelated with the second-stage error term, conditional on the control variables. Intuitively, distance to the closest US Ebola location should only affect cognitive performance in the interview through its effect on Ebola concern. While the exclusion restriction is not formally testable in the just-identified case, I report the results of two types of falsification tests in Table 2.3.

First, I check whether distance to the closest US Ebola case is associated with systematic regional differences in attitudes towards diseases or the tendency towards internet searches during epidemics. Such a correlation would cast doubt on the exclusion restriction because it seems plausible that attitudes towards diseases are related to cognitive function, as illustrated by the case of misperceptions. To test this, I check whether the instrument predicts Ebola search interest before the first US case and the level of concern about another recent epidemic, the swine flu pandemic of 2009/2010, in my main sample of October interviews. As demonstrated by the small and insignificant coefficient estimates in columns (1) and (2) of the table, there is no relationship between distance to US Ebola locations and internet searches for diseases before the emergency of Ebola cases in the US.

In a second, more direct falsification exercise, I test whether the instrument is related to cognitive function in a placebo sample of HRS interviews that should be unaffected by Ebola concern in October. In particular, I check whether there is a relationship between placebo variants of the instrument and the cognitive test scores of HRS respondents interviewed in September. Based on their interview dates, these respondents should be similar to October respondents, but they were interviewed before the emergence of the first US Ebola case. Therefore, I expect their cognitive function test scores to be unrelated to their distance to the eventual US Ebola locations if the exclusion restriction holds. I compute two different placebo versions of the proposed instrument that keep the locations of US cases unchanged but backdate their occurrence by different amounts of time. For the placebo instrument in column (3), I assume that all US Ebola cases happened before September 2014. In contrast, the placebo instrument used in column (4) assumes that all Ebola cases happened exactly 30 days before their actual dates. In both cases, the estimated coefficient is small and insignificant, suggesting that my proposed instrument does not simply pick up pre-existing differences in cognitive function that are unrelated to US Ebola cases.[21]

Overall, my falsification tests provide no indication for potential violations of the exclusion restriction.

2.5.3 IV Estimates

The main results of my IV analysis are shown in Table 2.4. Column (1) estimates the reduced form. Under the maintained assumption that the exclusion restriction holds, the reduced form provides a test of the effect of worry on cognitive function that is valid irrespective of any

21. In unreported regressions, I verify that this conclusion carries over to regressions on the sample of all HRS 2014 interviews conducted before the first US case.

Table 2.3. Distance to US Ebola Locations and Selected Outcomes: Falsification Tests

	Main sample		Placebo sample	
Dependent variable:	Ebola searches before first US case	Swine flu concern	Verbal reasoning score	
	(1)	(2)	(3)	(4)
Distance to closest Ebola location	−0.024 (0.406)	0.285 (0.336)		
Distance to closest Ebola location (placebo)			−0.120 (0.789)	
Distance to closest Ebola location (placebo 2)				0.241 (0.680)
Demographic controls	Yes	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes	Yes
Location controls	Yes	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes	Yes
Observations	495	495	569	569
R^2 (adjusted)	0.272	0.377	0.415	0.415
Mean (dependent variable)	27.372	40.760	504.942	504.942

Notes: OLS estimates for the main sample and a placebo sample of HRS 2014 respondents interviewed in September (excluding September 30), with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola searches before first US case* is the relative search interest for the term “Ebola” in the media market of the interview location in the first week of August. *Swine flu concern* is the relative search interest for the topic “swine flu (swine influenza)” in the media market of the interview location between April 25, 2009, and August 10, 2010. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. *Distance to closest Ebola location (placebo)* is the logarithm of the kilometer distance to the closest of the three cities that have a connection to one of the four Ebola case diagnosed in the US between September 30 and October 23, even though the interview date precedes these cases. *Distance to closest Ebola location (placebo 2)* is the logarithm of the kilometer distance to the closest location that would have a relationship to a publicly known US Ebola case on the day of the interview if all US Ebola cases happened exactly 30 days earlier. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case (except in column (1)), and media market–level relative search interest for the topics “anxiety” and “virus” in 2013. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

weak-instrument concerns because it does not rely on a strong first-stage relationship. In line with the hypothesis that worrying about an Ebola epidemic impairs cognitive function, test performance increases significantly with distance to the closest US Ebola location, conditional on the full set of controls. Specifically, a one-unit increase in logarithmic distance implies an increase in verbal reasoning score of about 3.0 points on the W scale.

Table 2.4. Ebola Concern and Cognitive Function: IV Estimates

	Dependent variable: Verbal reasoning score		
	Reduced form	IV	OLS
	(1)	(2)	(3)
Distance to closest Ebola location	2.983***		
	(1.127)		
Ebola concern		−0.323**	−0.286***
		(0.130)	(0.095)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	495	495	495
R^2 (adjusted)	0.436	0.441	0.441
Mean (dependent variable)	505.115	505.115	505.115
Effective F-statistic		12.563	
Weak-instrument-robust confidence set		[-0.720, -0.102]	

Notes: Reduced-form and IV estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. *Ebola concern* is the average weekly relative search interest for the term "Ebola" in the media market of the interview location over all weeks between September 28 and the week of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern*. Column (3) reproduces the OLS estimates from Table 2.1. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

The corresponding IV estimate is presented in column (2). The estimated coefficient implies that a one SD increase in Ebola concern causes a reduction in verbal reasoning score of 4.66 points on the W scale. This is equivalent to a reduction in the probability of answering a test item of equivalent difficulty correctly from 50 to 36.6 or 90 to 84 percent, respectively (Jaffe, 2009). Comparing the level of Ebola concern in October 2014 with the counterfactual situation in which public concern is equal to the pre-October search interest peak, the extrapolated effect of worrying about a US Ebola epidemic would even amount to a 12-point drop in cognitive test performance.

In the bottom of the table, I also report a weak-instrument-robust confidence set based on test inversion of the cluster-robust version of the Anderson-Rubin test, using a significance level of five percent (Anderson and Rubin, 1949; Chernozhukov and Hansen, 2008; Andrews, Stock, and Sun, 2019). Like the reduced form estimates, this procedure does not rely on a strong first stage, so the resulting confidence interval survives any concerns about weak instruments. Therefore, it is reassuring that the null hypothesis of no effect is clearly rejected also using this alternative inference approach.

Column (3) reproduces the OLS estimates from Table 2.1. The IV estimate is approximately 13 percent larger, suggesting that the OLS results cannot be explained by remaining endogeneity problems. On the contrary, any potential downward endogeneity bias in the OLS estimates seems to be outweighed by the effect of measurement error in Ebola concern. The findings of the IV analysis therefore corroborate the conclusion that worrying about an epidemic induces a substantial cognitive cost.

2.5.4 Robustness

In Appendix 2.C, I report estimates from various additional analyses designed to test the robustness of the OLS and IV estimation results. I briefly summarize these analyses here.

First, I verify that my findings are not an artifact of the way my measure of local Ebola concern on interview day is constructed from the Google Trends search interest data. The Ebola concern measure employed in the main text, which is based on the average weekly search interest across all weeks between the first US Ebola case and the interview week, assumes that worries outlast the period in which individuals actively search for information about Ebola on the internet. In Table 2.C.1, I rerun the analyses under the assumption that Ebola concern equals search interest in the week of the interview. A downside of this measure is that measurement error due to US Ebola cases after the interview day in the same week plays a larger role here. Consequently, the resulting OLS estimate is close to zero, whereas the first-stage and IV estimates are similar to the respective main text coefficients. I steer a middle course between both approaches in table Table 2.C.2, which contains estimates based on measuring Ebola concern as the average search interest in the interview week and the preceding week.[22] For this third measure of Ebola concern, all estimates closely resemble those reported in the main text. The insignificance of the specific way Ebola concern is measured is also reflected in the empirical distributions of the three measures within my sample (see Table 2.B.1). While the dispersion naturally decreases for measures that average search interest over a larger number of weeks, the sample means all approximate a value of 60.

22. For one respondent with an early October interview, Ebola search interest data for the preceding September week in the respective media market is missing in more than 49 iterations of data collection. As a result, this measure of Ebola concern suffers from the selection bias described in Section 2.3.2 for the affected respondent. I exclude this respondent from the analysis.

Second, I assess the implications of a different way of aggregating the data from individual Google Trends queries into weekly search interest in the first place. Instead of the median, I now use the average across all 100 iterations.[23] The coefficient estimates of this robustness check, displayed in Table 2.C.3, are almost identical to those reported in the main text.

Third, I rerun the estimations for the subsample of HRS respondents that are interviewed in metropolitan areas. This should further alleviate concerns that the probability of US Ebola cases is lower in rural locations, but it comes at the cost of losing 20 percent of the sample. The estimated IV coefficient, presented in Table 2.C.4, shrinks by one-third in size and is only marginally statistically significant, but still qualitatively in line with the main results.

Finally, I check whether my results are robust to the use of alternative instruments. In their analyses of daily Twitter activity about Ebola in October 2014, Campante, Depetris-Chauvín, and Durante (2020) show that Ebola concern reacts most strongly to the first three US Ebola cases in Dallas, Texas, and Cleveland, Ohio. In contrast, the last US Ebola case in the City of New York, New York, on October 23 does not provoke a strong reaction anymore. This is intuitively plausible because every Ebola case that does not cause a US epidemic is good news about the contagiosity of the virus and the ability of the authorities to contain it. It also implies that proximity to early US Ebola cases only could also be a suitable instrument for Ebola concern. Based on this observation, I report results from IV analyses that use distance to the first or the closest of the first two US Ebola locations as alternative instruments in Table 2.C.5 and Table 2.C.6. As expected, the alternative instruments have a stronger first stage. In particular, the effective F-statistic is well beyond all relevant critical values for distance to the first US Ebola case. The corresponding reduced-form and IV estimates are qualitatively in line with those reported in the main text, with estimated IV coefficients of −0.253 and −0.326. All estimates are significantly different from zero at least at the five percent level. This indicates that the results of the IV analysis do not depend on a specific way of formalizing the idea that proximity to recent US Ebola cases increases worry. Moreover, it dispels any remaining concerns about weak instruments.

In sum, the results of these robustness checks show that the main findings are not sensitive to specific data management or analysis choices.

2.6 Do Interview Characteristics Respond to US Ebola Cases?

One unaddressed threat to the internal validity of my findings arises because interview dates are partly dictated by the availability of HRS participants rather than predetermined. This opens up the possibility that interview characteristics are endogenous to local Ebola concern. In particular, respondents from locations close to US Ebola cases might postpone or cancel their interview for fear of infection. If individuals with positive unobservable shocks to cognitive

23. For media market–weeks where data is missing for some queries because of the Google Trends privacy threshold, I use the midpoint between a lower and an upper bound on the average, where the lower bound results from setting missing values to zero and the upper bound from setting missing values to the lowest observed value. This only concerns search interest before the first US case for a few main-sample respondents, not the main regressor of interest.

ability in 2014 are more likely to respond to nearby Ebola cases by delaying their interview, the documented negative impact of Ebola concern on cognitive test performance could also be explained by a selection effect. However, note that such a relationship seems implausible in light of the role of misperceptions about Ebola for the spread of fear, described in Section 2.2. Assuming that the likelihood of holding wrong beliefs about Ebola decreases in cognitive ability, respondents with positive shocks to cognitive functioning should be less likely to avoid an interview for fear of infection after the occurrence of a nearby Ebola case.

Nonetheless, I also assess the plausibility of selection due to differential avoidance behavior by cognitive ability empirically. As a first step, I look for indications that the emergence of nearby US Ebola cases affects interview characteristics at all. Table 2.5 presents estimates from regressions of various interview characteristics on distance to the closest US Ebola location, using the same set of covariates as the previous analyses. First, I check whether the likelihood of telephone interviews decreases with distance to US Ebola cases, as would be expected if respondents avoid face-to-face interviews for fear of getting infected. The respective coefficient estimate is reported in column (1) of the table. In contrast to the prediction, a one-unit increase in logarithmic distance is associated with an eight-percentage-point higher likelihood of a telephone interview. This estimate is significantly different from zero at the ten percent level.

Hesitant respondents should also require more call attempts by the interviewer until they agree to be interviewed. This prediction is tested in column (2) and finds support in the data. In particular, the number of interactions between interviewer and respondent until the interview is conducted decreases by about 1.4 calls for each unit increase in logarithmic distance, which is highly significant.

To check for delayed interviews, I extend the sample to all HRS 2014 respondents with interview dates after the first US Ebola case on September 30. Column (3) contains the coefficient estimate from a regression of interview date, measured in days since the first interview of the 2014 HRS, on distance to the closest Ebola location in this extended sample. On average, interviews are conducted about eight days earlier for a one-unit increase in logarithmic distance to US Ebola cases.

In Table 2.B.2, I show that there are no comparable differences in interview characteristics by distance to the three subsequent Ebola locations in the sample of all HRS interviews conducted before the first US case. Thus, the data does not provide an indication that the significant effects detected in Table 2.5 could be an artifact of location-specific data collection procedures like differences in the start dates of the interview period. Yet, the directions of the observed differences also do not unequivocally support the idea that they are driven by respondent reactions to Ebola cases, as this is difficult to reconcile with a decrease in telephone interviews in proximity to Ebola locations. Arguably, switching to a telephone interview would be the easiest way to eliminate the risk of an Ebola infection even for individuals who hold wrong beliefs about the transmission dynamics of the disease. In line with this reasoning, HRS employees expressed in private communication to me that they are not aware of protocol changes or respondent concerns about contracting Ebola during interviews conducted in 2014.

For my purposes, the relevant question is whether changes in interview characteristics by distance to Ebola locations can explain my results. In this respect, it is reassuring that I control

Table 2.5. Distance to US Ebola Locations and Interview Characteristics

	Main sample		All interviews after first US case		
Dependent variable:	Telephone interview	Number of call attempts	Interview date (days since wave start)		
	(1)	(2)	(3)	(4)	(5)
Distance to closest Ebola location	0.079*	−1.384***	−7.698***	−7.084**	−8.209**
	(0.043)	(0.497)	(1.952)	(3.129)	(3.278)
Distance to closest Ebola location × High predicted score				−1.362	
				(2.216)	
Distance to closest Ebola location × High predicted score (RF)					0.989
					(2.308)
Demographic controls	Yes	Yes	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes	Yes	Yes
Location controls	Yes	Yes	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes	Yes	Yes
Observations	495	495	1358	1358	1358
R^2 (adjusted)	0.052	0.105	0.356	0.357	0.355
Mean (dependent variable)	0.432	16.030	274.772	274.772	274.772

Notes: OLS estimates for two different samples of HRS 2014 participants who were interviewed after the first US Ebola case, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. In columns (4) and (5), I report block-bootstrapped standard errors that additionally account for the presence of generated regressors. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. *High predicted score* and *High predicted score (RF)* are indicators for respondents with an above-median predicted verbal reasoning score, based on their covariates and coefficient estimates of an OLS or random forest regression of verbal reasoning score on these covariates in the sample of all HRS 2014 respondents interviewed before the first US Ebola case. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview (except in column (1)), the number of call attempts until an interview was conducted (except in column (2)), and interview week dummies (except in columns (3) to (5)). Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. Columns (4) and (5) additionally include the main effect of *High predicted score* and *High predicted score (RF)*, respectively. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

for interview mode and the number of contact attempts in the main specifications. Delayed interviews in proximity to Ebola cases, on the other hand, would only change the interpretation of my results if they were positively correlated with respondents' cognitive capacity in 2014. Therefore, I next investigate whether the strength of the documented relationship between nearby Ebola cases and interview dates is related to observable determinants of cognitive function in the comprehensive HRS data.[24] To this end, I implement the following two-stage estimation procedure.

First, I model verbal reasoning score as a function of observable respondent characteristics in the sample of all HRS 2014 interviews before the first US Ebola case. The estimation equation is

$$VerbalReasoningScore_{jc} = \alpha + \boldsymbol{\eta}_1'\boldsymbol{X}_j + \boldsymbol{\eta}_2'\boldsymbol{C}_c + \boldsymbol{\eta}_3'\boldsymbol{M}_m + \rho_r + \epsilon_{jc}, \tag{2.2}$$

where respondents are denoted by the subscript j rather than i to indicate the use of a different sample. As predictors, I include the full set of demographic, baseline cognition, and location controls as well as census region dummies, but not interview characteristics, which might themselves depend on Ebola distance for interviews conducted after the first US case. The resulting model can explain 41.4 percent of the variation in actual cognitive test scores. Then, I use the estimated coefficients from this model to predict test scores in the sample of HRS interviews conducted after the first US Ebola case and generate an indicator variable $\widehat{HighPredictedScore}_{ic}$ that is equal to one for respondents with an above-median predicted test score in this sample. The result is an easily interpretable summary measure of counterfactual cognitive ability in the absence of US Ebola cases.

The second stage builds on the regression model underlying column (3) of Table 2.5, which I augment with the constructed indicator for high predicted test scores as a main and interaction effect with Ebola distance. The respective estimation equation is thus given by

$$\begin{aligned} InterviewDate_{ict} = {} & \alpha + \beta_1 EbolaDistance_{ct} + \beta_2 EbolaDistance_{ct} \times \widehat{HighPredictedScore}_{ic} \\ & + \kappa_0 \widehat{HighPredictedScore}_{ic} + \boldsymbol{\kappa}_1'\boldsymbol{X}_i + \boldsymbol{\kappa}_2'\boldsymbol{C}_c + \boldsymbol{\kappa}_3'\boldsymbol{M}_m + \rho_r + \epsilon_{ict}, \end{aligned} \tag{2.3}$$

where the coefficient on the interaction term, β_2, reveals whether selection into later interview dates differs by cognitive ability.[25] The resulting coefficient estimates of the second stage are reported in column (4) of the table. The main effect of Ebola distance, which corresponds to the average delay of the interview per unit of logarithmic distance for respondents with below-median predicted cognitive performance, is similar to the equivalent estimate for the whole

24. Note that I cannot simply use the realized verbal reasoning test score of October respondents because I expect this to itself be affected by distance to Ebola locations via the hypothesized effect of Ebola concern.

25. The variable $\widehat{HighPredictedScore}_{ic}$ is itself estimated and therefore subject to sampling uncertainty which is not accounted for in standard covariance estimates (Pagan, 1984). To correct the standard errors for the presence of a generated regressor and allow for arbitrary intra-cluster correlation within media markets, I implement a two-step block bootstrap approach that follows Ashraf and Galor (2013). Each bootstrap replicate is generated in the following way, resembling the two-stage estimation procedure: First, a random sample of HRS respondents is drawn with replacement from the sample of all HRS respondents with interview dates before the first US Ebola case. I run the prediction regression (Equation 2.2) on this random sample and save the resulting OLS coefficients. Second, I draw a random sample of media markets with replacement from the sample of all HRS respondents interviewed after the first US Ebola case, thereby accounting for clustering within media markets. I use the saved OLS coefficients from the prediction regression to predict the counterfactual verbal reasoning scores in this random sample of post-September HRS interviews and construct the indicator variable for above-median predicted test scores. I then estimate the second-stage OLS regression (Equation 2.3), which yields the bootstrap coefficient estimates of this replicate. This process of two-step block bootstrap sampling and OLS estimation is repeated 1000 times. The standard deviations of the resulting sample of 1000 bootstrap coefficient estimates from the second-stage regression are the bootstrap standard errors of the point estimates of these coefficients.

sample in column (3). The interaction effect, though noisily estimated, is close to zero. This suggests that potential selection effects are not related to cognitive ability.

I also implement a variant of the estimation procedure that replaces the linear OLS model in the initial prediction stage with a random forest regression model (Breiman, 2001; Hastie, Tibshirani, and Friedman, 2009). Random forests are a machine learning algorithm that is known to have a high predictive performance in a variety of contexts even with little tuning.[26] In contrast to OLS, it is a nonparametric method that also works well in settings characterized by higher-order interactions and nonlinearities, i.e., exactly when OLS would be expected to perform poorly. Moreover, the algorithm ensures that the support of the resulting predictions is bounded by the range of observed outcome values in the data the model is trained on. Therefore, it provides an ideal complement to the linear prediction model underlying the estimates in column (4).

As before, I use the random forest regression model to predict counterfactual verbal reasoning scores for all HRS respondents interviewed after the first US Ebola case, which are then converted into an indicator variable $\widehat{HighPredictedScore}_{ic}^{RF}$ by means of a median split to ease interpretation.[27] The model is trained on the sample of HRS 2014 interviews before the first US Ebola case, using the same pool of available covariates as the linear prediction model. I use 200 decision trees and do not restrict the depth of the individual trees. The quality of potential splits is measured by mean squared error. The number of observations required to split an internal tree node l and the fraction of randomly chosen covariates available for the algorithm to choose from for a given split m are determined by hyperparameter tuning, resulting in a choice of $(l, m) = (25, 0.3)$.

To assess the predictive performance of the tuned model, I use a nested cross-validation procedure with five folds each in both the inner tuning loop and the outer validation loop.[28] The resulting average out-of-sample R^2 of predicted continuous test scores across the five testing folds of the outer cross-validation loop is 0.405. For the dichotomized summary measure $\widehat{HighPredictedScore}_{ic}^{RF}$, 74.0 percent of respondents are classified correctly on average.

Column (5) of Table 2.5 displays the resulting second-stage estimates with block-bootstrapped standard errors. Relative to the variant of the estimation procedure with the linear prediction model in column (4), the conclusion does not change. If anything, the point

26. See Mullainathan and Spiess (2017) for a discussion of how machine learning algorithms can be used for prediction tasks in economics and Athey and Imbens (2019) for a general overview of machine learning methods for economists.

27. Alternatively, a random forest classifier could be used to predict directly whether respondents obtain an above-median test score. I choose the regression model because the additional information in continuous test scores reduces the likelihood of misclassifying people with extreme scores.

28. Especially with a large number of included predictors, machine learning algorithms are prone to overfitting: the model parameters may pick up specific random components of the sample at hand that are not present in other samples, implying an overestimation of their predictive performance on new data. The typical remedy is to randomly split the available data into a separate training and testing sample, where only the former is used to train the model and the latter to evaluate its performance. Cross-validation is a structured approach to model evaluation in which model training and evaluation are repeated several times and each observation appears in a testing sample exactly once. This is accomplished by splitting the available data into several folds, with each fold serving as the testing sample for a model trained on the remaining folds.

estimate of β_2^{RF} suggests a weakly positive interaction effect, implying a weaker association between Ebola distance and interview dates for respondents with high predicted cognitive test scores.

All in all, the empirical investigation provides no indication that changes in interview characteristics are related to cognitive function, suggesting that selection into interview dates is unlikely to explain my main results.

2.7 Discussion

The preceding sections establish a robust negative association between the local level of Ebola concern and performance in a test of fluid intelligence in a sample of HRS participants. The results of an IV strategy confirm the findings of the OLS analysis and lead to the conclusion that the documented relationship is causal. My investigations also show that this cannot be explained by a selection effect due to differential avoidance behavior by cognitive ability, but rather implies a direct influence of Ebola concern on cognitive test performance. Regarding the mechanisms behind these results, two points for discussion emerge from my analysis.

The first point pertains to the interpretation of my Ebola concern measures, which are constructed from data on internet search volume. While especially the reduced-form and IV estimates provide strong evidence that cognitive costs are ultimately caused by the emergence of nearby Ebola cases, it is less clear whether the effect is driven by emotional responses to these cases in the form of worry or fear or an increase in factual public interest.

There is reason to believe that Ebola concern captures emotional expressions. This fits well with evidence from nationally representative surveys that many people were in fact worried about Ebola and with anecdotal evidence of fear-based behavioral overreactions, summarized in Section 2.2. Moreover, the strong dependence of the level of Ebola concern on geographical distance to the closest case location—both within the US and with respect to Ebola cases in other countries—suggests that the perceived threat of infection plays an important role. In particular, it seems difficult to explain why the WHO's declaration of an international public health emergency after almost 2000 Ebola cases and 1000 deaths in Africa would be objectively less interesting than the occurrence of a single Ebola case in the US several weeks later, if not because of a lower perceived personal threat.[29] Yet, the nationwide level of Ebola concern is more than twice as high after the first US case than after the WHO announcement. The reading that internet activity reveals emotional reactions is also supported by van Lent, Sungur, Kunneman, van de Velde, and Das (2017), who carry out a content analysis of a random sample of tweets about Ebola that were posted in the Netherlands between March and November 2014. They find explicit expressions of fear in one-fifth of all analyzed tweets, and the fraction and number of fearful tweets increased substantially in October, when initial Ebola cases in nearby European countries became public. Therefore, this is my preferred interpretation.

To the extent that my measures of Ebola concern merely pick up variation in factual public interest, my results suggest that the resulting attentional capture nonetheless leads to similar

29. A similar argument could be made with respect to the difference in information value between an Ebola case in a neighboring and a more distant US state.

levels of mental preoccupation. This would be more in line with the scarcity hypothesis advanced by Mullainathan and Shafir (2013), which posits that the perception of pressing needs like a scarcity of income or time taxes cognitive capacity by inducing a narrow attentional focus on the perceived problem.

Alternatively, respondents might consciously choose to allocate less cognitive effort to the interview because nearby Ebola cases change the marginal benefit of other cognitively demanding activities. For instance, local levels of concern about Ebola might affect the utility of consuming and reflecting about the latest news. This would be in the spirit of Altmann, Grunewald, and Radbruch (Forthcoming), who show that relative economic incentives affect the allocation of attention among different choice domains to the detriment of neglected domains. However, this seems unlikely in this case because (i) cognitive effort contributes little to solving verbal analogies, and (ii) it is hard to imagine how a reduction in mental focus on the interview would benefit competing activities. In particular, the interview would not be noticeably shorter if respondents performed worse on the cognitive tests and fruitfully engaging in any other activity simultaneously with the cognitive tests seems close to impossible.

A second discussion point concerns the importance of contextual factors of the 2014 US Ebola episode for the size of the documented effect. As described in Section 2.2, the occurrence of Ebola cases in the US led to a surge in sensationalist media reporting about Ebola. At the same time, politicians strategically exploited the situation to appeal to voters in light of the upcoming 2014 midterm elections, likely contributing to the spread of fear (Campante, Depetris-Chauvín, and Durante, 2020). I cannot disentangle whether the cognitive costs of US Ebola cases are primarily attributable to concern about the perceived threat of the cases themselves or to its amplification by these other factors. This suggests caution with respect to the generalizability of my findings to other settings, as similar levels of epidemic activity will not necessarily induce similar levels of public concern.

2.8 Conclusion

My findings suggest the possibility of an unexplored economic cost of epidemics: negative emotional responses to the threat of the disease could cause suboptimal choices and reduced cognitive productivity by taking a toll on individuals' cognitive functioning. The cognitive cost of worrying about an epidemic could in principle affect everyone, even those who are unaffected by the disease itself or resulting government interventions, and in all domains of decision-making. However, those who have more to fear, e.g., because they work in occupations that require interpersonal contact, might be hit hardest.

For policymakers, this implies that it is important to not only contain the epidemic itself, but also to curb excessive fears that may arise from it. In particular, information campaigns can prevent misperceptions about the danger of the disease. On the personal level, it seems wise to foresee the cognitive consequences of epidemic-induced worry and restrain one's ambitions and plans accordingly. For instance, one could defer important life decisions until fear has subsided rather than taking the risk of a suboptimal choice during the onset of a frightening epidemic.

Future research should investigate the specific economic consequences of epidemic-induced cognitive costs, especially with respect to cognitive errors in individual decision-making and reductions in labor productivity. Understanding which domains and individuals are most affected would pave the way for the development of strategies and targeted interventions to mitigate any negative impact. Moreover, it would be interesting to learn more about the underlying mechanisms of the documented reductions in cognitive functioning. The hypothesized underlying cognitive processes imply a reallocation of scarce attentional resources towards the perceived threat. It remains an open question whether this also leads to measurably better health outcomes such as a lower risk of infection.

Appendix 2.A HRS Data Appendix

This appendix describes the source and construction of HRS variables that are not explained in detail in the main text. The main regression specifications include a wide range of control variables that are constructed from survey items in the RAND HRS Longitudinal File and the RAND HRS Fat Files for the 2010, 2012, and 2014 wave. HRS control variables can be categorized into three groups. The respective source items of each described variable are given in parentheses and refer to the RAND HRS Longitudinal File unless noted otherwise.

Demographic controls include age, age squared (both based on item R12AGEY_E), and dummies for gender, race, education, and changes in household status, employment status, and health since the previous HRS wave in 2012. I control for gender by including a dummy variable equal to one for female respondents (based on item RAGENDER). Race and ethnicity are captured by separate dummies for identifying as White Hispanic, Black, other race Hispanic, or other race non-Hispanic, such that White non-Hispanic is the omitted category (based on items RARACEM and RAHISPAN). To account for education, I include four dummies that mark respondents' highest attained education category (based on item RAEDUC): one for passing the General Educational Development test, one for completing at most high school, one for having done some college, and one for holding at least a college degree, with less than high school forming the base category. The remaining dummy variables indicate changes in personal circumstances since the last HRS wave in 2012 that could be associated with cognitive decline. In particular, I construct a dummies for having started to live alone (based on items H11HHRES and H12HHRES), having moved into a nursing home (based on items R11NHMLIV and R12NHMLIV), reporting to have retired from work (based on items R11RETEMP and R12RETEMP), and having been diagnosed with dementia (item R12DEMENS) or a stroke (item R12STROKS).[30] For each categorical variable, I also include a dummy equal to one if one of the underlying survey items is missing for the respondent to avoid losing observations due to missing control variables.

Baseline cognition controls are the scores of four different cognitive function tests conducted in the 2012 HRS wave and the changes in these scores between the 2010 and 2012 wave. In particular, I include the W score of a quantitative reasoning test (based on items R10NSSCRE and R11NSSCRE) and the number of correct answers in a test of retrieval fluency (where the 2012 score is constructed as the difference between items ND196 and ND198 in the RAND HRS 2012 Fat File, and the 2010 score as the difference between items MD196 and MD198 in the RAND HRS 2010 Fat File) to control for baseline differences in fluid intelligence across respondents.[31] The remaining test scores are the number of correct answers in a word recall test (based on items R10TR20 and R11TR20) and the number of correct subtractions in the serial sevens test (based on items R10SER7 and R11SER7), both of which are tests aimed at detecting early warning signs of cognitive decline. For the latter two tests, the

30. An equivalent dummy for having been diagnosed with Alzheimer's disease since the last wave (item R12ALZHES) is collinear with the other control variables in the main sample and therefore omitted.

31. For a small number of respondents, retrieval fluency scores are negative because the count of correct answers is exceeded by the number of incorrect answers. I allow for negative scores in my analyses. However, the results are very similar if negative scores are truncated at zero instead.

HRS imputes missing values based on non-changing baseline demographics, wave-specific demographic, economic, and health status variables, and previous- and current-wave cognitive function measures. I also use imputed values as control variables in my analysis because the imputation procedure relies on pre-Ebola characteristics only.

The last group of HRS control variables are interview controls. These consist of an indicator for telephone interviews, for which in-person interviews are the omitted category (item OB084 in the RAND HRS 2014 Fat File), and the number of call or contact attempts by the interviewer until an interview is conducted (item O085 in the RAND HRS 2014 Fat File). In addition, this group also includes interview week dummies constructed from restricted data on the beginning date of the interview. Interview weeks are defined to start on Sundays for consistency with the weekly search interest data provided by Google Trends, which follows the same convention.

Appendix 2.B Additional Tables

Table 2.B.1. Descriptive Statistics for the Main Sample

Variable	Mean	SD	Min.	Median	Max.	Obs.
Verbal reasoning score	505.115	29.888	<440	498	560	495
Ebola concern	63.786	14.438	(s)	(s)	(s)	495
Distance to closest Ebola location	6.813	0.903	(s)	(s)	(s)	495
Female	0.622	0.485	0	1	1	495
Age	66.038	10.162	<45	64	95	495
White Hispanic	0.077	0.266	0	0	1	495
Black	0.206	0.405	0	0	1	495
Other race Hispanic	0.022	0.148	0	0	1	495
Other race non-Hispanic	0.055	0.227	0	0	1	495
General Education Development test	0.053	0.223	0	0	1	495
High school degree	0.253	0.435	0	0	1	495
Some college	0.265	0.442	0	0	1	495
College degree	0.275	0.447	0	0	1	495
Started living alone (since 2012)	0.075	0.263	0	0	1	495
Moved into nursing home (since 2012)	0.012	0.110	0	0	1	495
Retired (since 2012)	0.075	0.263	0	0	1	495
Diagnosed with dementia (since 2012)	0.014	0.118	0	0	1	495
Dementia status missing	<0.010	<0.100	0	0	1	495
Diagnosed with a stroke (since 2012)	0.018	0.134	0	0	1	495
Stroke status missing	<0.010	<0.100	0	0	1	495
Quantitative reasoning score (2012)	518.911	32.565	409	519	584	495
Change in quantitative reasoning score (2010 to 2012)	21.119	35.026	<−65.0	17.8	>125.0	495
Retrieval fluency score (2012)	17.513	7.191	<1	17	>35	495
Change in retrieval fluency score (2010 to 2012)	0.578	6.991	<−14	0	>20	495
Word recall score (2012)	10.394	3.162	<4	10	>17	495
Change in word recall score (2010 to 2012)	−0.295	3.291	<−7	0	>8	495
Serial sevens score (2012)	3.634	1.562	0	4	5	495
Change in serial sevens score (2010 to 2012)	−0.032	1.473	<−3	0	>3	495
Telephone interview	0.432	0.496	0	0	1	495
Number of call attempts	16.030	10.490	3	14	>60	495
Interview day (within October 2014)	14.600	9.210	(s)	(s)	(s)	495
Distance to closest large city	3.689	1.560	(s)	(s)	(s)	495
Ebola searches before first US case	27.372	5.981	(s)	(s)	(s)	495
Search interest for “anxiety” (2013)	55.242	5.983	(s)	(s)	(s)	495
Search interest for “virus” (2013)	50.805	5.882	(s)	(s)	(s)	495
Swine flu concern (2009 and 2010)	40.760	5.044	(s)	(s)	(s)	495
Ebola concern (interview week only)	61.502	29.665	(s)	(s)	(s)	495
Ebola concern (last two weeks)	58.777	22.775	(s)	(s)	(s)	494
Ebola concern (average)	63.859	14.488	(s)	(s)	(s)	495

(*continued on next page*)

Table 2.B.1. Descriptive Statistics for the Main Sample (*continued from previous page*)

Variable	Mean	SD	Min.	Median	Max.	Obs.
Distance to first Ebola location	7.117	0.673	(s)	(s)	(s)	495
Distance to closest Ebola location (excl. New York)	6.883	0.747	(s)	(s)	(s)	495

Notes: Mean, standard deviation, minimum, median, maximum, and number of observations of important variables for the main estimation sample. All distance variables are in log-transformed kilometers. (s) indicates statistics which are suppressed to preclude the possibility of inferential identification of small geographical areas, in line with the HRS disclosure rules for analyses based on restricted data. Some statistics are censored at the request of the HRS to avoid disclosure of small cell sizes.

Table 2.B.2. Distance to US Ebola Locations and Interview Characteristics: Placebo Regressions

	Placebo sample		All interviews before first US case
Dependent variable:	Telephone interview	Number of call attempts	Interview date (days since wave start)
	(1)	(2)	(3)
Distance to closest Ebola location (placebo)	0.009	−0.015	0.360
	(0.026)	(0.485)	(1.041)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	569	569	12 079
R^2 (adjusted)	0.162	0.277	0.329
Mean (dependent variable)	0.308	13.844	78.379

Notes: OLS estimates for two different samples of HRS 2014 participants who were interviewed before the first US Ebola case, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. The placebo sample used in columns (1) and (2) is the sample of all HRS 2014 respondents interviewed in September (excluding September 30th), while the sample used in column (3) also includes respondents from earlier months. *Distance to closest Ebola location (placebo)* is the logarithm of the kilometer distance to the closest of the three cities that have a connection to one of the four Ebola case diagnosed in the US between September 30 and October 23, even though the interview date precedes these cases. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview (except in column (1)), the number of call attempts until an interview was conducted (except in column (2)), and interview week dummies (except in column (3)). Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Appendix 2.C Robustness

Table 2.C.1. Alternative Measures of Ebola Concern: Search Interest in the Interview Week

Dependent variable:	Ebola concern (interview week only)	Verbal reasoning score	
	First stage	IV	OLS
	(1)	(2)	(3)
Distance to closest Ebola location	−8.705***		
	(1.984)		
Ebola concern (interview week only)		−0.343**	−0.056
		(0.131)	(0.074)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	495	495	495
R^2 (adjusted)	0.861	0.416	0.431
Mean (dependent variable)	61.502	505.115	505.115
Effective F-statistic	19.245	19.245	
Weak-instrument-robust confidence set		[-0.669, -0.099]	

Notes: First-stage, IV, and OLS estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern (interview week only)* is the relative search interest for the term "Ebola" in the media market of the interview location in the week of the interview. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern (interview week only)*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 2.C.2. Alternative Measures of Ebola Concern: Search Interest in the Last Two Weeks

Dependent variable:	Ebola concern (last two weeks)	Verbal reasoning score	
	First stage	IV	OLS
	(1)	(2)	(3)
Distance to closest Ebola location	−8.488***		
	(1.355)		
Ebola concern (last two weeks)		−0.354***	−0.238**
		(0.130)	(0.099)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	494	494	494
R^2 (adjusted)	0.873	0.435	0.437
Mean (dependent variable)	58.777	505.132	505.132
Effective F-statistic	39.258	39.258	
Weak-instrument-robust confidence set		[-0.627, -0.102]	

Notes: First-stage, IV, and OLS estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern (last two weeks)* is the average weekly relative search interest for the term "Ebola" in the media market of the interview location over the week of the interview and the preceding week. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern (last two weeks)*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 2.C.3. Alternative Aggregation of Google Trends Data

Dependent variable:	Ebola concern (average)	Verbal reasoning score		
	First stage	Reduced form	IV	OLS
	(1)	(2)	(3)	(4)
Distance to closest Ebola location	−9.327***	2.965***		
	(2.586)	(1.119)		
Ebola concern (average)			−0.318**	−0.279***
			(0.127)	(0.093)
Demographic controls	Yes	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes	Yes
Location controls	Yes	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes	Yes
Observations	495	495	495	495
R^2 (adjusted)	0.659	0.436	0.441	0.441
Mean (dependent variable)	63.859	505.115	505.115	505.115
Effective F-statistic	13.011		13.011	
Weak-instrument-robust confidence set			[-0.706, -0.101]	

Notes: First-stage, reduced-form, IV, and OLS estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern (average)* is the average weekly relative search interest for the term "Ebola" in the media market of the interview location over all weeks between September 28 and the week of the interview, using the average rather than the median of 100 queries on the Google Trends website. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. All control variables based on Google Trends data use the average rather than the median of 100 queries on the Google Trends website. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern (average)*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 2.C.4. Estimates for the Subsample of Metropolitan Area Interviews

Dependent variable:	Ebola concern	Verbal reasoning score		
	First stage	Reduced form	IV	OLS
	(1)	(2)	(3)	(4)
Distance to closest Ebola location	−9.315***	1.937		
	(2.744)	(1.218)		
Ebola concern			−0.208*	−0.251***
			(0.120)	(0.091)
Demographic controls	Yes	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes	Yes
Location controls	Yes	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes	Yes
Observations	394	394	394	394
R^2 (adjusted)	0.643	0.463	0.469	0.470
Mean (dependent variable)	63.837	504.668	504.668	504.668
Effective F-statistic	11.526		11.526	
Weak-instrument-robust confidence set			[-0.507, 0.063]	

Notes: First-stage, reduced-form, IV, and OLS estimates for the subsample of HRS respondents who were interviewed in a metropolitan area in October 2014, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern* is the average weekly relative search interest for the term “Ebola” in the media market of the interview location over all weeks between September 28 and the week of the interview. *Distance to closest Ebola location* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics “anxiety” and “virus” in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 2.C.5. Alternative Instruments: Distance to the First US Ebola Location

Dependent variable:	Ebola concern	Verbal reasoning score	
	First stage	Reduced form	IV
	(1)	(2)	(3)
Distance to first Ebola location	−14.546***	4.740***	
	(1.326)	(1.501)	
Ebola concern			−0.326***
			(0.107)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	495	495	495
R^2 (adjusted)	0.759	0.438	0.441
Mean (dependent variable)	63.786	505.115	505.115
Effective F-statistic	120.365		120.365
Weak-instrument-robust confidence set			[-0.541, -0.128]

Notes: First-stage, reduced-form, and IV estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern* is the average weekly relative search interest for the term “Ebola” in the media market of the interview location over all weeks between September 28 and the week of the interview. *Distance to first Ebola location* is the logarithm of the kilometer distance to Dallas, TX, the location where the first US Ebola case was diagnosed on September 30. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics “anxiety” and “virus” in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 2.C.6. Alternative Instruments: Distance to the Closest of the First Two US Ebola Locations

Dependent variable:	Ebola concern	Verbal reasoning score	
	First stage	Reduced form	IV
	(1)	(2)	(3)
Distance to closest Ebola location (excl. New York)	−12.236***	3.091**	
	(1.982)	(1.386)	
Ebola concern			−0.253**
			(0.108)
Demographic controls	Yes	Yes	Yes
Baseline cognition controls	Yes	Yes	Yes
Interview controls	Yes	Yes	Yes
Location controls	Yes	Yes	Yes
Census region dummies	Yes	Yes	Yes
Observations	495	495	495
R^2 (adjusted)	0.710	0.434	0.441
Mean (dependent variable)	63.786	505.115	505.115
Effective F-statistic	38.096		38.096
Weak-instrument-robust confidence set			[-0.479, -0.035]

Notes: First-stage, reduced-form, and IV estimates for the main sample, with standard errors in parentheses. Standard errors are robust to heteroscedasticity and arbitrary intra-cluster correlation within media markets. *Ebola concern* is the average weekly relative search interest for the term "Ebola" in the media market of the interview location over all weeks between September 28 and the week of the interview. *Distance to closest Ebola location (excl. New York)* is the logarithm of the kilometer distance to the closest location with a relationship to a US Ebola case that was publicly known on the day of the interview, excluding the location of the last US Ebola case diagnosed on October 23. Demographic controls include age, age squared, and dummies for gender, race, education, and changes in household status (living alone, moving into a nursing home), employment status (retirement), and health (being diagnosed with dementia or a stroke) since the previous HRS wave in 2012. Baseline cognition controls are scores from tests of quantitative reasoning, retrieval fluency, word recall, and the serial sevens test in 2012 and changes in these scores between 2010 and 2012. Interview controls are an indicator for a telephone interview, the number of call attempts until an interview was conducted, and interview week dummies. Location controls include the logarithm of the kilometer distance to the closest large city, Ebola search interest before the first US Ebola case, and media market–level relative search interest for the topics "anxiety" and "virus" in 2013. The effective F-statistic for the first stage reported in the bottom of the table accounts for the use of cluster-robust standard errors. The weak-instrument-robust confidence set reported in the bottom of the table is constructed by inverting the cluster-robust version of an Anderson-Rubin test at the five percent level for the coefficient on *Ebola concern*. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

References

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh. 2020. "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys." *Journal of Public Economics* 189: 104245. DOI: 10.1016/j.jpubeco.2020.104245. [61]

Altmann, Steffen, Andreas Grunewald, and Jonas Radbruch. Forthcoming. "Interventions and Cognitive Spillovers." *Review of Economic Studies*, DOI: 10.1093/restud/rdab087. [86]

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *Annals of Mathematical Statistics* 20 (1): 46–63. DOI: 10.1214/aoms/1177730090. [79]

Andrews, Isaiah, James H. Stock, and Liyang Sun. 2019. "Weak Instruments in Instrumental Variables Regression: Theory and Practice." *Annual Review of Economics* 11: 727–53. DOI: 10.1146/annurev-economics-080218-025643. [75, 79]

Ashraf, Quamrul, and Oded Galor. 2013. "The 'Out of Africa' Hypothesis, Human Genetic Diversity, and Comparative Economic Development." *American Economic Review* 103 (1): 1–46. DOI: 10.1257/aer.103.1.1. [83]

Athey, Susan, and Guido W. Imbens. 2019. "Machine Learning Methods That Economists Should Know About." *Annual Review of Economics* 11: 685–725. DOI: 10.1146/annurev-economics-080217-053433. [84]

Bacher-Hicks, Andrew, Joshua Goodman, and Christine Mulhern. 2021. "Inequality in Household Adaptation to Schooling Shocks: Covid-Induced Online Learning Engagement in Real Time." *Journal of Public Economics* 193: 104345. DOI: 10.1016/j.jpubeco.2020.104345. [66]

Baker, Scott R., Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. 2020. "How Does Household Spending Respond to an Epidemic? Consumption during the 2020 COVID-19 Pandemic." *Review of Asset Pricing Studies* 10 (4): 834–62. DOI: 10.1093/rapstu/raaa009. [61]

Baker, Scott R., and Andrey Fradkin. 2017. "The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data." *Review of Economics and Statistics* 99 (5): 756–68. DOI: 10.1162/REST_a_00674. [66]

Balgova, Maria, Simon Trenkle, Christian Zimpelmann, and Nico Pestel. 2022. "Job Search during a Pandemic Recession: Survey Evidence from the Netherlands." *Labour Economics* 75: 102142. DOI: 10.1016/j.labeco.2022.102142. [60]

Banerjee, Abhijit V., and Esther Duflo. 2007. "The Economic Lives of the Poor." *Journal of Economic Perspectives* 21 (1): 141–67. DOI: 10.1257/jep.21.1.141. [62]

Bedrosian, Sara R., Cathy E. Young, Laura A. Smith, Joanne D. Cox, Craig Manning, Laura Pechta, Jana L. Telfer, Molly Gaines-McCollom, Kathy Harben, Wendy Holmes, Keri M. Lubell, Jennifer H. McQuiston, Kristen Nordlund, John O'Connor, Barbara S. Reynolds, Jessica A. Schindelar, Gene Shelley, and Katherine Lyon Daniel. 2016. "Lessons of Risk Communication and Health Promotion — West Africa and United States." *MMWR Supplements* 65 (3): 68–74. DOI: 10.15585/mmwr.su6503a10. [63]

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. 2013. "Who is 'Behavioral'? Cognitive Ability and Anomalous Preferences." *Journal of the European Economic Association* 11 (6): 1231–55. DOI: 10.1111/jeea.12055. [57]

Benson, Jamie, Stas Amato, Turner Osler, David Hosmer, and Ajai Malhotra. 2021. "Hard Lessons: The Impact of Secondary Education on COVID-19 Cases and Deaths in the U.S." Unpublished manuscript, University of Vermount. April 27, 2021. DOI: 10.2139/ssrn.3825107. [61]

Binder, Carola. 2020. "Coronavirus Fears and Macroeconomic Expectations." *Review of Economics and Statistics* 102 (4): 721–30. DOI: 10.1162/rest_a_00931. [60]

Bogliacino, Francesco, Cristiano Codagnone, Felipe Montealegre, Frans Folkvord, Camilo Gómez, Rafael Charris, Giovanni Liva, Francisco Lupiáñez-Villanueva, and Giuseppe A. Veltri. 2021. "Negative Shocks Predict Change in Cognitive Function and Preferences: Assessing the Negative Affect and Stress Hypothesis." *Scientific Reports* 11: 3546. DOI: 10.1038/s41598-021-83089-0. [61]

Breiman, Leo. 2001. “Random Forests.” *Machine Learning* 45 (1): 5–32. DOI: 10.1023/A:1010933404324. [84]

Brodeur, Abel, Andrew E. Clark, Sarah Fleche, and Nattavudh Powdthavee. 2021. “COVID-19, Lockdowns and Well-Being: Evidence from Google Trends.” *Journal of Public Economics* 193: 104346. DOI: 10.1016/j.jpubeco.2020.104346. [66]

Brodeur, Abel, David Gray, Anik Islam, and Suraiya Bhuiyan. 2021. “A Literature Review of the Economics of COVID-19.” *Journal of Economic Surveys* 35 (4): 1007–44. DOI: 10.1111/joes.12423. [60]

Brynjolfsson, Erik, John J. Horton, Adam Ozimek, Daniel Rock, Garima Sharma, and Hong-Yi TuYe. 2020. “COVID-19 and Remote Work: An Early Look at US Data.” NBER Working Paper No. 27344. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w27344. [60]

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini. 2009. “Cognitive Skills Affect Economic Preferences, Strategic Behavior, and Job Attachment.” *Proceedings of the National Academy of Sciences* 106 (19): 7745–50. DOI: 10.1073/pnas.0812360106. [57]

Campante, Filipe, Emilio Depetris-Chauvín, and Ruben Durante. 2020. “The Virus of Fear: The Political Impact of Ebola in the U.S.” NBER Working Paper No. 26897. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w26897. [59, 60, 64, 67, 72, 80, 86]

Caplin, Andrew, and John Leahy. 2001. “Psychological Expected Utility Theory and Anticipatory Feelings.” *Quarterly Journal of Economics* 116 (1): 55–79. DOI: 10.1162/003355301556347. [57]

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang. 2016. “Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday.” *American Economic Review* 106 (2): 260–84. DOI: 10.1257/aer.20140481. [61]

Chen, Haiqiang, Wenlan Qian, and Qiang Wen. 2021. “The Impact of the COVID-19 Pandemic on Consumption: Learning from High-Frequency Transaction Data.” *AEA Papers and Proceedings* 111: 307–11. DOI: 10.1257/pandp.20211003. [60]

Chernozhukov, Victor, and Christian Hansen. 2008. “The Reduced Form: A Simple Approach to Inference with Weak Instruments.” *Economics Letters* 100 (1): 68–71. DOI: 10.1016/j.econlet.2007.11.012. [79]

Chevalier, Michelle S., Wendy Chung, Jessica Smith, Lauren M. Weil, Sonya M. Hughes, Sibeso N. Joyner, Emily Hall, Divya Srinath, Julia Ritch, Prea Thathiah, Heidi Threadgill, Diana Cervantes, and David L. Lakey. 2014. “Ebola Virus Disease Cluster in the United States — Dallas County, Texas, 2014.” *Morbidity and Mortality Weekly Report* 63 (46): 1087–88. URL: https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6346a11.htm?s_cid=mm6346a11_w. [63]

Coltart, Cordelia E. M., Benjamin Lindsey, Isaac Ghinai, Anne M. Johnson, and David L. Heymann. 2017. “The Ebola Outbreak, 2013–2016: Old Lessons for New Epidemics.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 372 (1721): 20160297. DOI: 10.1098/rstb.2016.0297. [62]

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman. 2020. “Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data.” *Brookings Papers on Economic Activity* 51 (2): 35–82. DOI: 10.1353/eca.2020.0006. [60]

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. “In Search of Attention.” *Journal of Finance* 66 (5): 1461–99. DOI: 10.1111/j.1540-6261.2011.01679.x. [66]

de Bruijn, Ernst-Jan, and Gerrit Antonides. 2022. “Poverty and Economic Decision Making: A Review of Scarcity Theory.” *Theory and Decision* 92 (1): 5–37. DOI: 10.1007/s11238-021-09802-7. [62]

Dean, Emma Boswell, Frank Schilbach, and Heather Schofield. 2019. “Poverty and Cognitive Function.” In *The Economics of Poverty Traps*. Edited by Christopher B. Barrett, Michael R. Carter, and Jean-Paul Chavas. Chicago, IL: University of Chicago Press. Chapter 2, 57–118. DOI: 10.7208/9780226574448-004. [61]

Dean, Joshua T. 2021. “Noise, Cognitive Function, and Worker Productivity.” Unpublished manuscript, University of Chicago Booth School of Business. March 30, 2021. URL: https://joshuatdean.com/wp-content/uploads/2021/08/NoiseCognitiveFunctionandWorkerProductivity.pdf. [70]

Deck, Cary, and Salar Jahedi. 2015. "The Effect of Cognitive Load on Economic Decision Making: A Survey and New Experiments." *European Economic Review* 78: 97–119. DOI: 10.1016/j.euroecorev.2015.05.004. [58]

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. 2010. "Are Risk Aversion and Impatience Related to Cognitive Ability?" *American Economic Review* 100 (3): 1238–60. DOI: 10.1257/aer.100.3.1238. [57]

Duquennois, Claire. 2022. "Fictional Money, Real Costs: Impacts of Financial Salience on Disadvantaged Students." *American Economic Review* 112 (3): 798–826. DOI: 10.1257/aer.20201661. [61]

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2021. "The Macroeconomics of Epidemics." *Review of Financial Studies* 34 (11): 5149–87. DOI: 10.3386/w26882. [60]

Elster, Jon. 1998. "Emotions and Economic Theory." *Journal of Economic Literature* 36 (1): 47–74. [57, 60]

Eysenck, Michael W., Nazanin Derakshan, Rita Santos, and Manuel G. Calvo. 2007. "Anxiety and Cognitive Performance: Attentional Control Theory." *Emotion* 7 (2): 336–53. DOI: 10.1037/1528-3542.7.2.336. [57]

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2021. "Internal and External Effects of Social Distancing in a Pandemic." *Journal of Economic Theory* 196: 105293. DOI: 10.1016/j.jet.2021.105293. [60]

Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth. 2021. "Coronavirus Perceptions and Economic Anxiety." *Review of Economics and Statistics* 103 (5): 968–78. DOI: 10.1162/rest_a_00946. [60]

Fisher, Gwenith G., John J. McArdle, Ryan J. McCammon, Amanda Sonnega, and David R. Weir. 2014. "New Measures of Fluid Intelligence in the HRS: Quantitative Reasoning, Verbal Reasoning, Verbal Fluency." HRS Documentation Report DR-027. Ann Arbor, MI: Survey Research Center, Institute for Social Research, University of Michigan. URL: https://hrs.isr.umich.edu/sites/default/files/biblio/dr-027b.pdf. [65]

Flückiger, Matthias, Markus Ludwig, and Ali Sina Önder. 2019. "Ebola and State Legitimacy." *Economic Journal* 129 (621): 2064–89. DOI: 10.1111/ecoj.12638. [60]

Goeijenbier, Marco, Jeroen J. A. van Kampen, Chantal B. E. M. Reusken, Mario P. G. Koopmans, and Eric C. M. van Gorp. 2014. "Ebola Virus Disease: A Review on Epidemiology, Symptoms, Treatment and Pathogenesis." *Netherlands Journal of Medicine* 72 (9): 442–48. URL: https://www.njmonline.nl/getpdf.php?id=1495. [62]

Gomes, Marcelo F. C., Ana Pastore y Piontti, Luca Rossi, Dennis Chao, Ira Longini, M. Elizabeth Halloran, and Alessandro Vespignani. 2014. "Assessing the International Spreading Risk Associated with the 2014 West African Ebola Outbreak." *PLoS Currents* 6: DOI: 10.1371/currents.outbreaks.cd818f63d40e24aef769dda7df9e0da5. [63]

González-Torres, Ada, and Elena Esposito. 2020. "Epidemics and Conflict: Evidence from the Ebola Outbreak in Western Africa." Unpublished manuscript, Ben-Gurion University of the Negev, Department of Economics. December 19, 2020. DOI: 10.2139/ssrn.3544606. [60]

Goolsbee, Austan, and Chad Syverson. 2021. "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020." *Journal of Public Economics* 193: 104311. DOI: 10.1016/j.jpubeco.2020.104311. [57, 60]

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd edition. New York, NY: Springer. DOI: 10.1007/b94608. [84]

Haushofer, Johannes, and Ernst Fehr. 2014. "On the Psychology of Poverty." *Science* 344 (6186): 862–67. DOI: 10.1126/science.1232491. [61]

Health and Retirement Study. 2019a. "Cross-Wave Geographic Information (Detail) 1992–2016 - Version 7.1 (Early) restricted dataset [dataset]." Ann Arbor, MI: University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) [producer and distributor]. [65]

Health and Retirement Study. 2019b. "Respondent Date of Interview 1992–2016 - Version 7.0 (Early) restricted dataset [dataset]." Ann Arbor, MI: University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) [producer and distributor]. [65]

Heckman, James J., Jora Stixrud, and Sergio Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *Journal of Labor Economics* 24 (3): 411–82. DOI: 10.1086/504455. [57]

Jacob, Brian, and Jesse Rothstein. 2016. "The Measurement of Student Ability in Modern Assessment Systems." *Journal of Economic Perspectives* 30 (3): 85–108. DOI: 10.1257/jep.30.3.85. [66]

Jaffe, Lynne E. 2009. "Development, Interpretation, and Application of the W Score and the Relative Proficiency Index." Woodcock-Johnson III Assessment Service Bulletin No. 11. Rolling Meadows, IL: Riverside Publishing. URL: https://www.hmhco.com/$%5Csim$/media/sites/home/hmh-assessments/clinical/woodcock-johnson/pdf/wjiii/wj3_asb_11.pdf. [66, 71, 78]

Kaufman, Bruce E. 1999. "Emotional Arousal as a Source of Bounded Rationality." *Journal of Economic Behavior & Organization* 38 (2): 135–44. DOI: 10.1016/s0167-2681(99)00002-5. [58]

Kostova, Deliana, Cynthia H. Cassell, John T. Redd, Desmond E. Williams, Tushar Singh, Lise D. Martel, and Rebecca E. Bunnell. 2019. "Long-Distance Effects of Epidemics: Assessing the Link between the 2014 West Africa Ebola Outbreak and U.S. Exports and Employment." *Health Economics* 28 (11): 1248–61. DOI: 10.1002/hec.3938. [60]

Kremer, Michael, Gautam Rao, and Frank Schilbach. 2019. "Behavioral Development Economics." In *Handbook of Behavioral Economics: Foundations and Applications 2*. Edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson. Amsterdam: North-Holland. Chapter 5, 345–458. DOI: 10.1016/bs.hesbe.2018.12.002. [61]

Lichand, Guilherme, and Anandi Mani. 2020. "Cognitive Droughts." Department of Economics Working Paper Series No. 341. Zurich: University of Zurich. URL: https://www.econ.uzh.ch/static/wp/econwp341.pdf. [61]

Loewenstein, George. 2000. "Emotions in Economic Theory and Economic Behavior." *American Economic Review* 90 (2): 426–32. DOI: 10.1257/aer.90.2.426. [57, 60]

Loewenstein, George F., Elke U. Weber, Christopfer K. Hsee, and Ned Welch. 2001. "Risk as Feelings." *Psychological Bulletin* 127 (2): 267–86. DOI: 10.1037/0033-2909.127.2.267. [60]

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science* 341 (6149): 976–80. DOI: 10.1126/science.1238041. [61, 65, 71]

Mathews, Andrew. 1990. "Why Worry? The Cognitive Function of Anxiety." *Behaviour Research and Therapy* 28 (6): 455–68. DOI: 10.1016/0005-7967(90)90132-3. [57]

McCarty, Carolyn L., Colin Basler, Mateusz Karwowski, Marguerite Erme, Gene Nixon, Chris Kippes, Terry Allan, Toinette Parrilla, Mary DiOrio, Sietske de Fijter, Nimalie D. Stone, David A. Yost, Susan A. Lippold, Joanna J. Regan, Margaret A. Honein, Barbara Knust, and Christopher Braden. 2014. "Response to Importation of a Case of Ebola Virus Disease — Ohio, October 2014." *Morbidity and Mortality Weekly Report* 63 (46): 1089–91. URL: https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6346a12.htm?s_cid=mm6346a12_w. [63]

Montiel Olea, José Luis, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics* 31 (3): 358–69. DOI: 10.1080/00401706.2013.806694. [74, 75]

Moran, Tim P. 2016. "Anxiety and Working Memory Capacity: A Meta-Analysis and Narrative Review." *Psychological Bulletin* 142 (8): 831–64. DOI: 10.1037/bul0000051. [57]

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means So Much*. New York, NY: Henry Holt and Company. [58, 61, 86]

Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives* 31 (2): 87–106. DOI: 10.1257/jep.31.2.87. [84]

Pagan, Adrian. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25 (1): 221–47. DOI: 10.2307/2648877. [83]

Park, R. Jisung. 2022. "Hot Temperature and High Stakes Performance." *Journal of Human Resources* 57 (2): 400–34. DOI: 10.3368/jhr.57.2.0618-9535R3. [70]

Rainisch, Gabriel, Jason Asher, Dylan George, Matt Clay, Theresa L. Smith, Christine Kosmos, Manjunath Shankar, Michael L. Washington, Manoj Gambhir, Charisma Atkins, Richard Hatchett, Tim Lant, and Martin I. Meltzer. 2015. "Estimating Ebola Treatment Needs, United States." *Emerging Infectious Diseases* 21 (7): 1273–75. DOI: 10.3201/eid2107.150286. [63]

RAND Center for the Study of Aging. 2020. “RAND HRS Longitudinal File 2016 (V2) [dataset].” Santa Monica, CA: RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration [producer]. Ann Arbor, MI: University of Michigan [distributor]. Available at https://hrsdata.isr.umich.edu. [65]

RAND Center for the Study of Aging. 2021a. “RAND HRS 2010 Fat File (V5F) [dataset].” Santa Monica, CA: RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration [producer]. Ann Arbor, MI: University of Michigan [distributor]. Available at https://hrsdata.isr.umich.edu. [65]

RAND Center for the Study of Aging. 2021b. “RAND HRS 2012 Fat File (V3A) [dataset].” Santa Monica, CA: RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration [producer]. Ann Arbor, MI: University of Michigan [distributor]. Available at https://hrsdata.isr.umich.edu. [65]

RAND Center for the Study of Aging. 2021c. “RAND HRS 2014 Fat File (V2B) [dataset].” Santa Monica, CA: RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration [producer]. Ann Arbor, MI: University of Michigan [distributor]. Available at https://hrsdata.isr.umich.edu. [65]

Rassy, Dunia, and Richard D. Smith. 2013. “The Economic Impact of H1N1 on Mexico’s Tourist and Pork Sectors.” *Health Economics* 22 (7): 824–34. DOI: 10.1002/hec.2862. [60]

Robinson, Oliver J., Katherine Vytal, Brian R. Cornwell, and Christian Grillon. 2013. “The Impact of Anxiety upon Cognition: Perspectives from Human Threat of Shock Studies.” *Frontiers in Human Neuroscience* 7: 203. DOI: 10.3389/fnhum.2013.00203. [57]

Sari, Berna A., Ernst H. W. Koster, and Nazanin Derakshan. 2017. “The Effects of Active Worrying on Working Memory Capacity.” *Cognition and Emotion* 31 (5): 995–1003. DOI: 10.1080/02699931.2016.1170668. [57]

Sood, Gaurav. 2016. “Geographic Information on Designated Media Markets (V9) [dataset].” Harvard Dataverse [distributor]. Available at https://doi.org/10.7910/DVN/IVXEHT. [67]

Staiger, Douglas, and James H. Stock. 1997. “Instrumental Variables Regression with Weak Instruments.” *Econometrica* 65 (3): 557–86. DOI: 10.2307/2171753. [75]

SteelFisher, Gillian K., Robert J. Blendon, and Narayani Lasala-Blanco. 2015. “Ebola in the United States — Public Reactions and Implications.” *New England Journal of Medicine* 373 (9): 789–91. DOI: 10.1056/nejmp1506290. [63, 64, 67]

Stock, James H., and Motohiro Yogo. 2005. “Testing for Weak Instruments in Linear IV Regression.” In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*. Edited by Donald W. K. Andrews and James H. Stock. Cambridge: Cambridge University Press. Chapter 5, 80–108. DOI: 10.1017/CBO9780511614491.006. [75]

Towers, Sherry, Shehzad Afzal, Gilbert Bernal, Nadya Bliss, Shala Brown, Baltazar Espinoza, Jasmine Jackson, Julia Judson-Garcia, Maryam Khan, Michael Lin, Robert Mamada, Victor M. Moreno, Fereshteh Nazari, Kamaldeen Okuneye, Mary L. Ross, Claudia Rodriguez, Jan Medlock, David Ebert, and Carlos Castillo-Chavez. 2015. “Mass Media and the Contagion of Fear: The Case of Ebola in America.” *PLoS ONE* 10 (6): e0129179. DOI: 10.1371/journal.pone.0129179. [64]

U.S. Census Bureau. 2014. “National Places Gazetteer Files: 2014 [online].” Available at https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.2014.html. Accessed January 7, 2021. [68]

U.S. Census Bureau. 2019. “Centers of Population by County: 2010 [online].” Available at https://www.census.gov/geographies/reference-files/2010/geo/2010-centers-population.html. Accessed April 14, 2020. [68]

U.S. Census Bureau. 2020. “Annual Estimates of the Resident Population for Incorporated Places in the United States: April 1, 2010 to July 1, 2019 (SUB-IP-EST2019-ANNRES) [online].” Available at https://www.census.

gov/data/tables/time-series/demo/popest/2010s-total-cities-and-towns.html. Accessed January 7, 2021. [68]

van Lent, Liza G. G., Hande Sungur, Florian A. Kunneman, Bob van de Velde, and Enny Das. 2017. "Too Far to Care? Measuring Public Attention and Fear for Ebola Using Twitter." *Journal of Medical Internet Research* 19 (6): e193. DOI: 10.2196/jmir.7219. [85]

Vincenty, Thaddeus. 1975. "Direct and Inverse Solutions of Geodesics on the Ellipsoid with Application of Nested Equations." *Survey Review* 23 (176): 88–93. DOI: 10.1179/sre.1975.23.176.88. [68]

WHO Ebola Response Team. 2014. "Ebola Virus Disease in West Africa — The First 9 Months of the Epidemic and Forward Projections." *New England Journal of Medicine* 371 (16): 1481–95. DOI: 10.1056/nejmoa1411100. [62]

Witters, Dan, and Jim Harters. 2020. "Worry and Stress Fuel Record Drop in U.S. Life Satisfaction [online]." Gallup News. Available at https://news.gallup.com/poll/310250/worry-stress-fuel-record-drop-life-satisfaction.aspx. Accessed May 17, 2021. [57]

World Health Organization. 2016. "Ebola Virus Disease Situation Report – 10 June 2016." URL: https://apps.who.int/iris/handle/10665/208883. [62]

Yacisin, Kari, Sharon Balter, Annie Fine, Don Weiss, Joel Ackelsberg, David Prezant, Ross Wilson, David Starr, Jennifer Rakeman, Marisa Raphael, Celia Quinn, Amita Toprani, Nancy Clark, Nathan Link, Demetre Daskalakis, Aletha Maybank, Marcelle Layton, and Jay K. Varma. 2015. "Ebola Virus Disease in a Humanitarian Aid Worker — New York City, October 2014." *Morbidity and Mortality Weekly Report* 64 (12): 321–23. URL: https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6412a3.htm?s_cid=mm6412a3_w. [63]

Chapter 3

The Cost of Worrying about an Epidemic: Experimental Evidence on Labor Productivity during the COVID-19 Pandemic*

3.1 Introduction

A characteristic feature of human productivity is that it is responsive to short-term changes in psychological state. Consequently, productivity can tumble during emotionally challenging times. For instance, existing evidence points to productivity declines for individuals who experience personal tragedies like the unexpected death of a child (van den Berg, Lundborg, and Vikström, 2017) or the diagnosis of a severe disease (Liu and Zhu, 2014) as well as shocks to housing wealth during a recession (Bernstein, McQuade, and Townsend, 2021). Given the primary role of productivity in the economic system, changes in productivity can have far-reaching economic consequences. Productivity is linked to cross-country differences in economic prosperity, affects firm profitability and survival in market competition, and determines individuals' labor market outcomes (e.g., Hall and Jones, 1999; Foster, Haltiwanger, and Syverson, 2008; Lindqvist and Vestman, 2011). Therefore, preserving labor productivity during times of crisis is of crucial interest to policymakers, CEOs, and individuals alike.

This chapter asks whether labor productivity is adversely affected by emotional responses to the COVID-19 pandemic, one of the greatest public health crises in history and a major source of emotional distress. This question has direct implications for the optimal structure of productivity-preserving work environments and optimal public communication during epidemics. A negative causal effect would suggest that continuous exposure to information about the danger of the disease comes at the cost of reduced economic productivity, implying a trade-

* I thank Lorenz Goette, Andrew Oswald, and Florian Zimmermann for valuable comments and discussions, Markus Antony for administrative support related to project funding, and Holger Gerhardt for help with running the experiment at the BonnEconLab. I gratefully acknowledge funding by the Institute on Behavior and Inequality (briq) and the Bonn Graduate School of Economics (BGSE). The experiment in this chapter was preregistered in the AEA RCT Registry under identifying number AEARCTR-0006558.

off between the promotion of protective behavior and the intensity of the economic disruptions caused by the epidemic.

My main hypothesis is that worrying about the health consequences of an epidemic reduces labor productivity. It is based on studies from cognitive psychology, which indicate that worry can divert scarce attentional resources towards the perceived threat, thus reducing the available cognitive capacity for threat-unrelated processing tasks (Mathews, 1990; Eysenck, Derakshan, Santos, and Calvo, 2007; Robinson, Vytal, Cornwell, and Grillon, 2013; Moran, 2016; Sari, Koster, and Derakshan, 2017). Using field data from tests of cognitive ability during an episode of heightened public concern about Ebola in the US, Apenbrink (2021) provides evidence that this also applies to worrying about epidemic activity. Since most work tasks rely on cognitive skills, I expect the resulting short-term reductions in available cognitive function to impair labor productivity.[1] On the other hand, changes in cognitive function do not have to translate into changes of productivity. In contrast to many pure tests of cognitive ability, most real-life tasks also have an effort component. Thus, workers might be able to respond to a worry-induced reduction in available cognitive resources with a compensatory increase in mental effort.

To test the hypothesis, I conducted an online experiment with a student sample in autumn 2020, during the second wave of the COVID-19 pandemic in Germany. In the experiment, subjects can work on a cognitively demanding task that entails solving as many mental arithmetic problems as possible for a piece-rate wage. The task is divided into multiple five-minute blocks across two sessions on consecutive days, thus allowing to measure each subject's productivity and emotional state before and after introducing exogenous variation in emotional responses to COVID-19. Mental arithmetic resembles typical white-collar jobs in its reliance on both cognitive function and cognitive effort, making it an ideal task for the purposes of my experiment. Moreover, subject performance on tasks involving mental arithmetic is found to be responsive to induced changes in emotions in previous experimental studies (e.g., Oswald, Proto, and Sgroi, 2015). The main treatment manipulation exposes participants to real media reports selected to either amplify or alleviate their worries about COVID-19. A comparison of subjects' number of correct answers in the subsequent blocks of the cognitive task across the two experimental conditions thus identifies the causal effect of worrying about COVID-19 on labor productivity.

My first finding is that the selected media reports provoke a strong multi-faceted emotional response. Relative to the other experimental condition, subjects' average reported level of worry about the health consequences of COVID-19 for themselves and their loved ones is about 0.7 scale points higher after watching the worry-amplifying news video, conditional on worry on the previous day. The worry response is accompanied by an even greater relative reduction in happiness of about 1.4 points on an 11-point scale. These results show that my treatment manipulation successfully introduces exogenous variation in subjects' emotional reactions, consistent with the conjecture that the COVID-19 pandemic is a major source of emotional distress.

However, the induced emotional responses do not translate into economically meaningful differences in labor productivity. My preferred estimates indicate that exposure to the worry-

1. Existing research in economics suggests that concomitant other negative emotional responses like decreases in happiness could reduce productivity by means of similar cognitive processes (Kaufman, 1999; Oswald, Proto, and Sgroi, 2015).

evoking news video reduces the number of correct answers on the cognitive task by 0.137, which is a decrease of about 1.1 percent relative to the mean of the condition with the worry-alleviating media report. The estimated treatment effect is not significantly different from zero and precise enough to comfortably rule out an effect size of 0.2 standard deviations (SD) at the five percent significance level. This conclusion is confirmed when testing for distributional effects ($p = 0.712$ in a Kolmogorov-Smirnov randomization test) or an association between changes in productivity and changes in emotional state. Moreover, I do not find treatment effects on secondary productivity outcomes connected to specific cognitive processes that underlie productivity on the task, like the error rate or the average number of seconds subjects spend thinking on a mental arithmetic problem before submitting an answer. There are also no meaningful differences in posttreatment productivity in subgroups that partition the sample by participants' general tendency to worry or reported COVID-19 media exposure, two characteristics that might indicate stronger responsiveness to the treatment manipulation.

I consider the possibility that an effect of worry on labor productivity may be masked or crowded out by unintended ancillary effects that arise from the treatment or the experimental design but are unrelated to worrying about COVID-19. First, I investigate whether the worry-amplifying news video motivates subjects to work harder by shifting their beliefs about the financial consequences of the pandemic, thus offsetting the effect of lower cognitive function. Similar effects are observed in a recent field experiment by Kaur, Mullainathan, Oh, and Schilbach (2021). I find that the treatment has no effect on subjects' concerns about their financial situation or job prospects and on their self-reported productivity goal for the cognitive task, thus effectively ruling out this alternative explanation. Second, I assess whether my conclusions might be affected by the possibility of cheating on the cognitive task, e.g., in the form of using a voice-controlled calculator. However, a high observed error rate, a negative relationship between improvements in answer speed and accuracy, and a better within-subject performance on relatively easier mental arithmetic problems suggest that widespread cheating is equally implausible.

Beyond identifying the general effect of epidemic-induced worrying on productivity, my experimental design also allows an exploratory test of whether the high salience of the pandemic in the media exacerbates this effect. In the last two task blocks of the second session, headlines and lead paragraphs of current news articles are displayed to subjects for a few seconds at random intervals between mental arithmetic problems. A secondary cross-randomized treatment manipulation varies whether subjects see headlines about COVID-19 or about unrelated neutral topics. Using a difference-in-differences approach, I can investigate whether the effect of worry on productivity changes when the pandemic is saliently featured in the news, while accounting for productivity effects of the displayed news headlines that are unrelated to worrying about COVID-19. As expected given the initial null result, the estimated coefficients are again close to zero and statistically insignificant.

Taking stock, my analyses provide no indication that negative emotional responses to the COVID-19 pandemic impair labor productivity. The intuition underlying my hypothesis suggests two potential explanations for this result: either worrying does not impede cognitive ability in my setting, or subjects compensate for worry-induced cognitive effects by increasing their mental effort, in line with the notion of income targeting (e.g., Fehr and Goette, 2007; Duong, Chu, and Yao, Forthcoming).

Two sets of additional analyses provide some suggestive evidence against the first and in favor of the latter explanation. First, by tracking participants' interest in the headlines displayed to them during the cognitive task, my experimental setting allows to test for changes in cognition directly. Comparing the relative propensity to save news articles about COVID-19 between subjects who were exposed to the worry-amplifying and worry-alleviating media report, I find evidence consistent with worry-induced attentional capture. Second, I analyze whether the main treatment manipulation affects subjects' self-assessments of their productivity. My estimates indicate that subjects are significantly more likely to report distracting thoughts after watching the worry-amplifying news video, consistent with the idea that subjects feel less productive because they have to exert more effort to accomplish the cognitive task. However, the estimated treatment effects on three other indicators of lower perceived productivity are not or only marginally significant. While I do not view these two pieces of evidence as conclusive, they provide some indication in favor of worry-induced changes in cognition that subjects make up for by means of increased cognitive effort. Therefore, this is my preferred interpretation of the results.

With respect to the external validity of my findings, it is important to keep in mind the timing of the experiment. In autumn 2020, when the experiment was conducted, eight months had passed since the start of the pandemic. Therefore, my results can only speak to the effects of emotional responses after individuals have had time to adapt to the situation rather than during the onset of an epidemic, which presumably comes with higher degrees of uncertainty and scope for stronger emotional turmoil.

My study belongs to a rapidly emerging literature on the economic effects of the COVID-19 pandemic, reviewed by Brodeur, Gray, Islam, and Bhuiyan (2021). Important focus areas of existing studies are the impact of government restrictions to contain the pandemic on various economic outcomes (e.g., Coibion, Gorodnichenko, and Weber, 2020; Sheridan, Andersen, Hansen, and Johannesen, 2020) and the supply and demand shocks caused by voluntary reductions in consumption and labor supply by agents who fear contracting the virus (e.g., Brynjolfsson, Horton, Ozimek, Rock, Sharma, et al., 2020; Chetty, Friedman, Hendren, Stepner, and The Opportunity Insights Team, 2020; Cox, Ganong, Noel, Vavra, Wong, et al., 2020; Eichenbaum, Rebelo, and Trabandt, 2021; Goolsbee and Syverson, 2021). Presumably due to limited data availability, there is less work on the effects of the pandemic on economic productivity. A number of studies investigate the impact of COVID-19 on academic research productivity on the basis of surveys and the quantity of working paper publications (e.g., Kruger, Maturana, and Nickerson, 2020; Barber, Jiang, Morse, Puri, Tookes, et al., 2021). Calculating changes in productivity from survey responses of UK firms, Bloom, Bunn, Mizen, Smietanka, and Thwaites (2020) estimate decreases in within-firm productivity that are partially offset by a between-firm increase in average productivity driven by output contractions in low-productivity sectors. However, the extent to which self-assessments and purely quantity-based measures of research output capture all relevant dimensions of productivity is unclear. The few studies with access to more objective metrics—either based on internal company analytics tools or on benchmarking moves of elite chess players against a powerful chess engine—focus on the effects of pandemic-induced shifts towards remote working, which seem to impair productivity in cognitively demanding occupations while improving it for call center workers (Emanuel and Harrington, 2021; Gibbs, Mengel, and Siemroth, 2021; Künn, Seel, and Zegners, 2022). My experiment also features an objective measure of productivity, but it differs from the aforementioned studies by isolating the labor productivity consequences of workers' psychological responses to the pandemic.[2] Methodologically, my use of a controlled experiment complements recent field studies that document a drop in output as a consequence of an epidemic without allowing to cleanly pin down the underlying causal channels (e.g., Correia, Luck, and Verner, 2020). Thereby, my findings contribute to an understanding of the specific mechanisms by which an epidemic disrupts (or does not disrupt) the economy, which is a prerequisite for the development of measures to effectively alleviate its adverse economic effects.

My study also relates to a literature on the economic effects of media exposure (e.g., DellaVigna and La Ferrara, 2015). The media plays an important role in times of crisis. By providing accurate, readily understandable, and timely information to a broad audience, it can keep the public informed about the newest developments and counteract the spread of misperceptions. As documented by Bursztyn, Rao, Roth, and Yanagizawa-Drott (2021) and Simonov, Sacher, Dubé, and Biswas (2022) in the context of the COVID-19 pandemic, media coverage can shape individuals' behavior and thereby affect the spread of the disease. On the other hand, both the framing and the sheer volume of media reporting can also propagate worry and fear. In text analyses of the media coverage of the COVID-19 pandemic in the US, Sacerdote, Sehgal, and Cook (2020) show that articles by major US news outlets are overwhelmingly negative in tone and choice of covered topics, irrespective of real-world epidemiological developments. Since my main experimental manipulation relies on real media reports that differ in both information content and framing, it poses a direct test of the effect of differences in media reporting. Consequently, my results speak to potential negative side effects of fear-evoking media reports during times of crises. My null result on productivity suggests that, holding labor supply fixed, reductions in economic productivity as a consequence of worry-inducing media reports are unlikely, at least when people have had time to get used to the situation. In contrast, the estimated negative effect on happiness implies a direct welfare cost of exposure to media about the pandemic, in line with arguments developed in Zeckhauser (1996) and recommendations of government agencies.[3] Moreover, my findings on worry-induced changes in cognition suggest that disease-related worry is one factor driving the demand for pandemic-related news.[4] This could be an explanation for the predominantly pessimistic media coverage of the COVID-19 pandemic documented by Sacerdote, Sehgal, and Cook (2020).

2. Barber et al. (2021) document a correlation between health concerns and self-perceived productivity decline, but cannot establish a causal relationship because their data lack exogenous variation and the pandemic simultaneously affects several other determinants of productivity as well.

3. On their webpage about coping with pandemic-induced stress, the US Centers for Disease Control and Prevention caution against constant consumption of news about the COVID-19 pandemic: https://www.cdc.gov/coronavirus/2019-ncov/daily-life-coping/managing-stress-anxiety.html.

4. For other studies on determinants of the demand for news during the COVID-19 pandemic, see Castriota, Delmastro, and Tonin (2020) and Faia, Fuster, Pezone, and Zafar (Forthcoming).

Finally, I contribute to research on the effect of emotions on cognitive function and labor productivity. My experiment builds on and extends the results of Apenbrink (2021), who provides field evidence that worrying about the possibility of an epidemic impedes cognitive function more generally. I investigate the effect of worry in a different context—i.e., a different disease at a different point in the life cycle of epidemic activity, in a different sample, and with a different identification strategy—and focus on a specific economic outcome that might be affected by worry about epidemics. My findings qualify the economic significance of the cognitive effects of epidemic-induced emotions, implying that they either require levels of worry that are beyond the scope of an experimental manipulation or do not necessarily translate into reductions in labor productivity. As discussed in more detail in Section 3.5, this resonates with the results of related studies from behavioral development economics on the cognitive implications of financial strain, which seem to be more robust in settings with high uncertainty where strong worries are presumably more prevalent (e.g., Mani, Mullainathan, Shafir, and Zhao, 2013; Carvalho, Meier, and Wang, 2016; Lichand and Mani, 2020; Kaur, Mullainathan, et al., 2021). At the same time, my results contrast with the positive association between happiness and labor productivity documented by Oswald, Proto, and Sgroi (2015) in a similar experimental setup, again pointing to the importance of subtle situational factors. In particular, my findings suggest the possibility that the effect of changes in happiness on productivity might be stronger for increases than decreases or depend on the salience of a personal productivity goal.

The remainder of the chapter is structured as follows. The details of the experimental design are described in Section 3.2, followed by an outline of the preregistered hypotheses and evaluation methods in Section 3.3. The results of the experiment are presented in Section 3.4. In Section 3.5, I discuss unaddressed threats to the internal validity of the study and how my findings can be reconciled with existing work on the effect of emotions on productivity. Finally, Section 3.6 concludes.

3.2 Experimental Design

An analysis of the effect of worrying about an epidemic on labor productivity requires (i) a clean measure of productivity on a task that is susceptible to cognitive distractions and (ii) exogenous variation in the level of worry. Importantly, the actual level of the epidemic has to be kept constant to avoid confounding the effect of the emotional response with the effect of the epidemic itself. For instance, an epidemic could also affect labor productivity by triggering avoidance behavior, changing opportunity costs of working, or due to direct health effects of the disease. Satisfying these requirements with observational data is difficult. Instead, I employ an online experiment in the context of the COVID-19 pandemic.

The online experiment consists of two sessions on two consecutive days. In each session, participants can work on a cognitively demanding task. The task is divided into multiple consecutive five-minute blocks during which participants are supposed to solve as many mental arithmetic problems as they can for a piece-rate wage. The purpose of the first session is to obtain baseline measures of participants' productivity on the task as well as their self-reported level of worry about COVID-19 and momentary happiness. In the second session, I introduce exogenous variation in emotional responses to COVID-19 by showing subjects real media reports. The main treatment manipulation exposes subjects to a short news video clip about COVID-19 that is selected to either amplify (in condition *Worry-Amplifying*) or alleviate (in condition *Worry-Alleviating*) their worries about the pandemic. This is followed immediately by a second elicitation of worry and happiness that serves as a manipulation check. Afterwards, the mental arithmetic blocks of the second session begin. A cross-randomized second treatment manipulation is implemented while subjects are working on the cognitive task. During the last two blocks of arithmetic problems, subjects are exposed to headlines of current news articles that are either about COVID-19 (in subcondition *COVID-19 Headlines*) or about unrelated neutral topics (in subcondition *Neutral Headlines*) for a few seconds in between tasks. A comparison of subjects' performance on the task across the experimental conditions identifies whether fear-inducing media reports about the pandemic causally affect labor productivity and whether this effect is exacerbated by a high salience of the topic in the news.

The following subsections provide a detailed description of all components of the experimental design.

3.2.1 Cognitive Task

Participants' primary task throughout the experiment is to solve mental arithmetic problems. Each problem consists of adding and subtracting four two-digit numbers and includes one subtraction and two summation operations, as in the following example: $65 + 11 - 37 + 29$. The task is structured into blocks of five minutes, with short breaks between consecutive blocks. During each block, participants can work on the problems at their own pace and receive €0.10 for each correctly solved problem. The sequence of mental arithmetic problems within each block is the same for all subjects, and each block contains some easy and some more difficult problems by design.

This task is well-suited for the purposes of my study for a number of reasons. First, it relies heavily on a fundamental and ubiquitous cognitive function. Specifically, research from cognitive psychology suggests that solving mental arithmetic problems with multiple digits requires the interaction of all components of working memory (DeStefano and LeFevre, 2004). At the same time, working memory is an important ingredient to all real-life tasks that require individuals to at least temporarily store and manipulate information in mind. Second, and in contrast to many pure tests of cognitive ability, the mental arithmetic task also rewards cognitive effort. Such a combined requirement for cognitive ability and cognitive effort is typical for white-collar jobs and should therefore increase the generalizability to other cognitively demanding work tasks. Finally, it establishes comparability with previous experimental work on labor productivity, in which similar mental arithmetic tasks have been used. In particular, Oswald, Proto, and Sgroi (2015) show that exogenous increases in positive affect increase productivity on a mental arithmetic task, thus providing direct evidence that performance on the task is malleable in the context of changes in emotions.

One concern with implementing a mental arithmetic task—or any other cognitively demanding task—in an online experiment is that participants might cheat, e.g., by using a calculator. While it is not possible to fully prevent this,[5] the course of events for each arithmetic problem is designed to impede the most obvious forms of cheating: Each problem starts with a blank screen. After one, two, or three seconds, a blurred text appears. Participants can then display the unblurred arithmetic problem by simultaneously holding down the keys “Q” and “Enter” on their keyboard, thus preventing them from using their hands for other activities. As soon as they stop holding down “Q” and “Enter” or press any other key, the arithmetic problem is replaced by an input field, and participants have five seconds to type in their answer.[6]

Throughout the experiment, there are six blocks of mental arithmetic problems. Subjects complete two blocks of problems in the first session, which is identical in all experimental conditions. In these blocks, one mental arithmetic problem follows the other without interruption, in accordance with the above sequence of events. The purpose of the first block is to give subjects some experience with the task, whereas performance on the second block serves as a baseline measure of productivity. Of the four blocks in the second session, the first two, which I label “non-news blocks”, are identical in structure to those in the first session. By contrast, the last two blocks have a slightly different structure that is explained in Section 3.2.3, along with the secondary treatment manipulation these blocks are designed to enable.

The primary outcome of interest is the average number of correctly solved mental arithmetic problems in the two non-news blocks of the second session. This is a measure of subjects’ labor productivity after the main treatment manipulation. In addition, the structure of the task permits the measurement of three secondary productivity outcomes. The aim is to shed light on the effect of worry on specific cognitive processes that underlie subjects’ productivity on the cognitive task.[7] First, the task is designed to reward attention. To complete as many problems as possible, subjects have to unblur each newly appearing problem without delay. Being less attentive has a cost in the form of wasted time, with the consequence that fewer problems can be attempted within a given block. I measure attention by computing an average reaction time, defined as the average number of seconds until participants press “Q” and “Enter” after the display of the blurred arithmetic problem, where the average is taken across all attempted problems during non-news blocks. Second, the task allows to construct two measures of working memory. As a measure for the time participants take for the mental calculation of their answers, I average the number of seconds the unblurred problem is displayed across all problems they attempt during non-news blocks. Moreover, the quality of answers is measured by the fraction of attempted problems in these blocks that are answered incorrectly. However, it is important to realize that the two measures of working memory are not independent because of the natural trade-off between accuracy and speed.

5. For instance, subjects will always be able to work together in pairs, even though the instructions explicitly mandate working alone and the resulting effective hourly wage from pairing up would be quite low. I discuss indications of cheating in my data and whether this could plausibly affect my results in Section 3.5.1.

6. Screenshots of the experimental interface for the cognitive task are provided in Figure 3.E.1 in the appendix.

7. Some of the secondary outcomes depend on the number of attempted mental arithmetic problems within one or multiple blocks. For this purpose, I define attempted problems as problems that are completed entirely within the five-minute time limit of the block, i.e., I discard data from problems during which the block time limit runs out before the input field for the answer disappears.

3.2.2 Main Experimental Conditions: *Worry-Amplifying* and *Worry-Alleviating*

To induce exogenous variation in the level of worry about the pandemic, I use real media reports about COVID-19 that differ in framing and information content. This provides a natural way to manipulate worry that combines the provision of factual information with differences in conveyed feelings and atmosphere. Compared to standard priming interventions, this also aims to provoke stronger emotional responses by shifting worry both upward and downward rather than only drawing attention to pre-existing concerns. At the same time, it also keeps the overall salience of the topic constant, thus permitting a manipulation check and a separate analysis of the additional effect of salience.

In the practical implementation, all subjects are randomly assigned to one of two main experimental conditions, with randomization conducted on the subject level. The assignment determines which of two short news video clips subjects watch at the beginning of the second session, before they start working on that session's four blocks of mental arithmetic problems. In particular, the clips are selected to either amplify or alleviate subjects' worries about the pandemic. The initial selection was conducted on the basis of an extensive screening of the media reporting about COVID-19 in Germany. The chosen videos were then tested in a small pilot experiment to ensure that they induce changes in worry in the study population.

The fear-evoking video is displayed to subjects in the *Worry-Amplifying* condition. It was produced by German public television broadcaster Bayerischer Rundfunk for their program *Frankenschau aktuell* on August 19, 2020, and deals with potential long-term health effects of a COVID-19 infection. It features the story of a 51-year-old man who barely survived a severe COVID-19 infection and suffered permanent neurological damage as a result of the disease. As a consequence, he now has to relearn basic everyday tasks like speaking, sitting, and walking even though he did not have any preconditions before the infection. This is conveyed vividly during a short interview with the patient, during which he struggles to articulate his words in a comprehensible manner and has to be subtitled. The video also includes comments by doctors who talk about the unpredictability of COVID-19 infections and the wide range of different long-term effects that have been observed, including heart, lung, and liver damage as well as neurological disorders. At the end, it is reported that about 10 percent of COVID-19 infections take a severe course.[8]

In contrast, subjects in the *Worry-Alleviating* condition are exposed to the fear-mitigating video, which is taken from the German public television news program *ZDF heute* from November 20, 2020. It is a report about the possibility that a COVID-19 vaccine could be approved for use in Europe within a few weeks. Moreover, it informs about the start of preparations for vaccination centers designed to enable a fast roll-out of vaccinations once a vaccine is approved. The video strikes a more optimistic tone and conveys the hope that a remedy for the pandemic is on the horizon.

A comparison of subjects' subsequent performance on the cognitive task across the two conditions identifies whether and how emotional responses to the pandemic causally affect labor

8. I cannot make the videos publicly available due to copyright issues. However, they are available to interested readers upon request.

productivity. To encourage participants to pay attention, I specify an unconditional payment of €4.00 for watching the video and announce that there will be an attention check in the form of a question about the video content afterwards. This appeals to subjects' reciprocity and allows screening out those who are inattentive nonetheless.

3.2.3 Secondary Experimental Conditions: *COVID-19 Headlines* and *Neutral Headlines*

To test whether the effect of worrying on labor productivity is exacerbated by the salience of the pandemic, I implement a cross-randomized secondary treatment manipulation. Whereas the main treatment manipulation induces exogenous variation in the level of fear about COVID-19, here the aim is to induce variation in the salience of the topic for a given level of worry. The underlying idea is that the effect of worrying about COVID-19 might linger longer if people frequently encounter the topic again, thus keeping the emotion top-of-mind.

To accommodate this in the experiment, I slightly adapt the structure of the cognitive task in the last part of the second session and expose subjects to additional media reports while working. This allows for an exploratory test without reducing statistical power for the main treatment manipulation. After subjects have completed the non-news blocks of mental arithmetic problems, which deliver the outcome data for the main treatment manipulation, the course of events for the individual mental arithmetic problems changes in the remaining two blocks of the task. In these blocks, which I call "news blocks", headlines and lead paragraphs of current online newspaper articles are displayed to subjects at random intervals between tasks for a period of 10 seconds. The net working time is unaffected by the number of headlines and remains at five minutes, so that productivity in news blocks can be measured analogously to productivity in the preceding blocks.

The secondary treatment manipulation determines what kind of articles subjects are exposed to while working. Within each of the two main experimental conditions, subjects are randomized into one of two secondary experimental conditions, *COVID-19 Headlines* or *Neutral Headlines*, with equal probability. Subjects in the *COVID-19 Headlines* subcondition see headlines of articles about COVID-19. In contrast, subjects in the *Neutral Headlines* subcondition see headlines of articles from a wide range of neutral topics like science, sports, entertainment, politics, and economics. Importantly, all articles in this condition are unrelated to the COVID-19 pandemic. These are articles that people would encounter if there was no pandemic or if it was less prominently featured in the news. This choice of control group ensures that differences in productivity across secondary experimental conditions for a given realization of the main treatment manipulation are driven by the topic of article headlines. Any general distraction effect of being exposed to news headlines while working, which could be specific to one's emotional state, applies equally to both secondary conditions.

All articles are collected at around 5:00 in the morning on the day of the second session, so most of the headlines concern the previous day. To compile potential articles, I use a news aggregator of popular German news websites, including Bild.de, Die Welt, FAZ.net, Handelsblatt, n-tv.de, Spiegel Online, Süddeutsche.de, and t-online.de. All subjects within a secondary condition that participate in the experiment on the same day see the same articles at the same point in the sequence of arithmetic problems. On average, subjects are shown an article headline every second time they have entered an answer for a mental arithmetic problem. The random component in the display of news headlines implies that they cannot avoid paying attention to the headlines if they want to complete as many arithmetic problems as possible within the time limit.

3.2.4 Elicitation of Additional Outcomes

To complement the productivity measures and allow a suggestive exploration of the underlying mechanisms of a potential treatment effect, I gather some additional outcomes throughout the experiment.

Measures of worry and happiness

I elicit measures of two emotions through which the main treatment manipulation might reduce productivity: worry and happiness. In particular, subjects are asked for self-reports of their level of worry about the health consequences of COVID-19 for themselves and their loved ones as well as their level of momentary happiness. Both measures are elicited using a Likert scale response format with an 11-point scale from 0 to 10. They are collected twice throughout the experiment, once at baseline at the beginning of the first session, and once just after the main treatment manipulation. Since the items are presented as numbers with two defined endpoints and many intermediate values, I interpret the resulting responses as approximating interval data. The collected data permits a manipulation check to assess whether the main treatment manipulation indeed increases subjects' degree of worry about the pandemic. Moreover, it reveals whether potential treatment effects should be interpreted as purely driven by worry or as the joint effect of a multi-faceted emotional response.

News interest

While my main interest is in the reduced-form effect of emotions on labor productivity, my experimental design also permits a direct test of the hypothesized underlying changes in cognition. A common perception regarding the cognition effects of worry in both psychology and economics is that scarce cognitive resources are redirected towards the object of concern, thus reducing their availability for other tasks (e.g. Eysenck et al., 2007; Mullainathan and Shafir, 2013; Lichand and Mani, 2020).

To allow for a test of this idea in the context of my experiment, I collect a behavioral measure of news interest. Subjects can click on the article headlines displayed to them during news blocks of the cognitive task to save them at no cost. Only saved articles are displayed again at the end of the experiment, so that subjects can read them afterwards. The underlying idea is that headlines about COVID-19 may capture subjects' attention even though they want to remain concentrated for the upcoming mental arithmetic problem, and this attentional shift is revealed in the action of saving the article for later reading. Since, for various reasons, instances of attentional capture may be too weak to induce subjects to save an article, I interpret this as a measure of how interested subjects are in the respective news topic relative to their focus on the task. This is a stronger notion of cognitive resource reallocation than a simple shift in attention. Of course, saving an article can also be interpreted simply as a demand for news about the respective topic.

Perceived productivity

I also elicit subjects' satisfaction with their performance on the task and their perceptions of their own productivity. Satisfaction with own performance in the second session is elicited on an 11-point scale. To get a sense of perceived productivity, subjects are asked to guess whether they have solved more, about the same, or fewer tasks in the second session relative to baseline. Guesses are elicited separately for non-news blocks and news blocks. The resulting data can be used to check whether subjects underestimate or overestimate the effect of media reports on productivity.

Beliefs about COVID-19

The postexperimental questionnaire elicits beliefs about COVID-19 that might be affected by the news videos. In particular, I elicit probabilistic beliefs about the risk of infection, the risk of long-term health effects conditional on infection, and the risk of death conditional on infection for inhabitants of Germany within the next six months. In addition, I also ask subjects to provide a best guess of the number of months until a vaccine becomes widely available in Germany.

Other outcomes

Finally, I elicit a few other outcomes designed to assess the plausibility of alternative explanations or permit heterogeneity analyses. After the main treatment manipulation, but before subjects start working, I ask them for their goal number of correct answers for the second session. In the postexperimental questionnaire, I elicit two measures of worry about the personal financial consequences of the pandemic, also on 11-point scales, and two items from validated survey measures of positive and negative reciprocity developed by Falk, Becker, Dohmen, Huffman, and Sunde (2016). Moreover, I include the German version of the Penn State Worry Questionnaire (PSWQ), which measures an excessive tendency to worry using 16 items (Meyer, Miller, Metzger, and Borkovec, 1990; Glöckner-Rist and Rist, 2014).

3.2.5 Procedures

The experiment was programmed in oTree (Chen, Schonger, and Wickens, 2016) and conducted online with subjects from the pool of the BonnEconLab. A total of 305 subjects participated in three iterations of the experiment that were spread out over a period of two weeks in late November and early December of 2020. At this time, Germany was in the midst of the second wave of the pandemic and experienced a continuously high number of daily new COVID-19 cases despite containment measures. The first vaccine manufacturers had just published encouraging results from clinical trials, but no vaccine had been authorized for use yet in the European Union and the US, so there remained some uncertainty about the start of vaccination campaigns. However, eight months had passed since the declaration of a national epidemic in Germany, implying that people had had ample time to adapt to the existence of the virus.

The subject pool of the BonnEconLab mainly consists of university students.[9] While students face a low risk of suffering a severe course of COVID-19 because of their young age, surveys indicate that they had similar levels of worry about the pandemic at the time of the experiment. The reason is that their rationally lower propensity to worry about their own health is offset by more pronounced concerns about family and friends. Computing the fraction of respondents who report high or very high levels of worry in at least one of these two dimensions in a nationally representative German survey at the end of November (Presse- und Informationsamt der Bundesregierung, Berlin, 2021), I find that 63.5 percent of respondents below 30 years of age are classified as worried, compared to 58.4 percent of respondents aged 30 or older. Taking into account the importance of worry about others in the subject pool, my study explicitly concerns worry about the health consequences of COVID-19 both for subjects' themselves and for their loved ones.

The chronology of the main components of the experimental design is visualized in Figure 3.1. Subjects were invited for a two-session online experiment using hroot (Bock, Baetge, and Nicklisch, 2014). On the day of the first session, registered participants were sent a unique link by e-mail early in the morning and had time until midnight to complete the first part of the experiment. The first session started with the elicitation of subjects' sociodemographic characteristics and baseline measures of worry about COVID-19 and momentary happiness. Afterwards, participants received instructions about the cognitive task.[10] To check their understanding, they had to answer a set of comprehension questions. The questions also tested whether their device satisfied the technical requirements for the task.[11] As the last component of the first session, subjects completed two blocks of mental arithmetic problems.

Procedures for the second session were similar. Again, subjects received their unique link early in the morning and could start the session at their own pace, as long as they finished until the end of the day. First, they read the instructions for the second session and completed comprehension questions. Second, they watched the news video clip of their experimental condition and answered the attention-check question about it, followed by the manipulation check. Then, they completed the two non-news blocks of arithmetic problems, followed by an elicitation of perceived productivity for these two blocks. Afterwards, they worked on the two news blocks of the cognitive task, during which they were presented with news headlines according to their experimental condition. Finally, they completed a short questionnaire.

Since treatment-induced worry might fade over time, an important aim of the experimental procedures and instructions for the second session is to minimize the time gap between

9. See Snowberg and Yariv (2021) for a recent study on the external validity of economic experiments with student populations relative to representative samples. A main finding is that comparative statics are very similar across different participant pools.

10. An English translation of the instructions for both sessions of the experiment is provided in Appendix 3.D.

11. Subjects could not proceed to the next page of the experiment until all comprehension questions were answered correctly. If they made a mistake, they were told to reread the instructions and try again, without receiving feedback which questions were answered incorrectly. Thus, the probability that participants passed the comprehension questions by trial and error is very low.

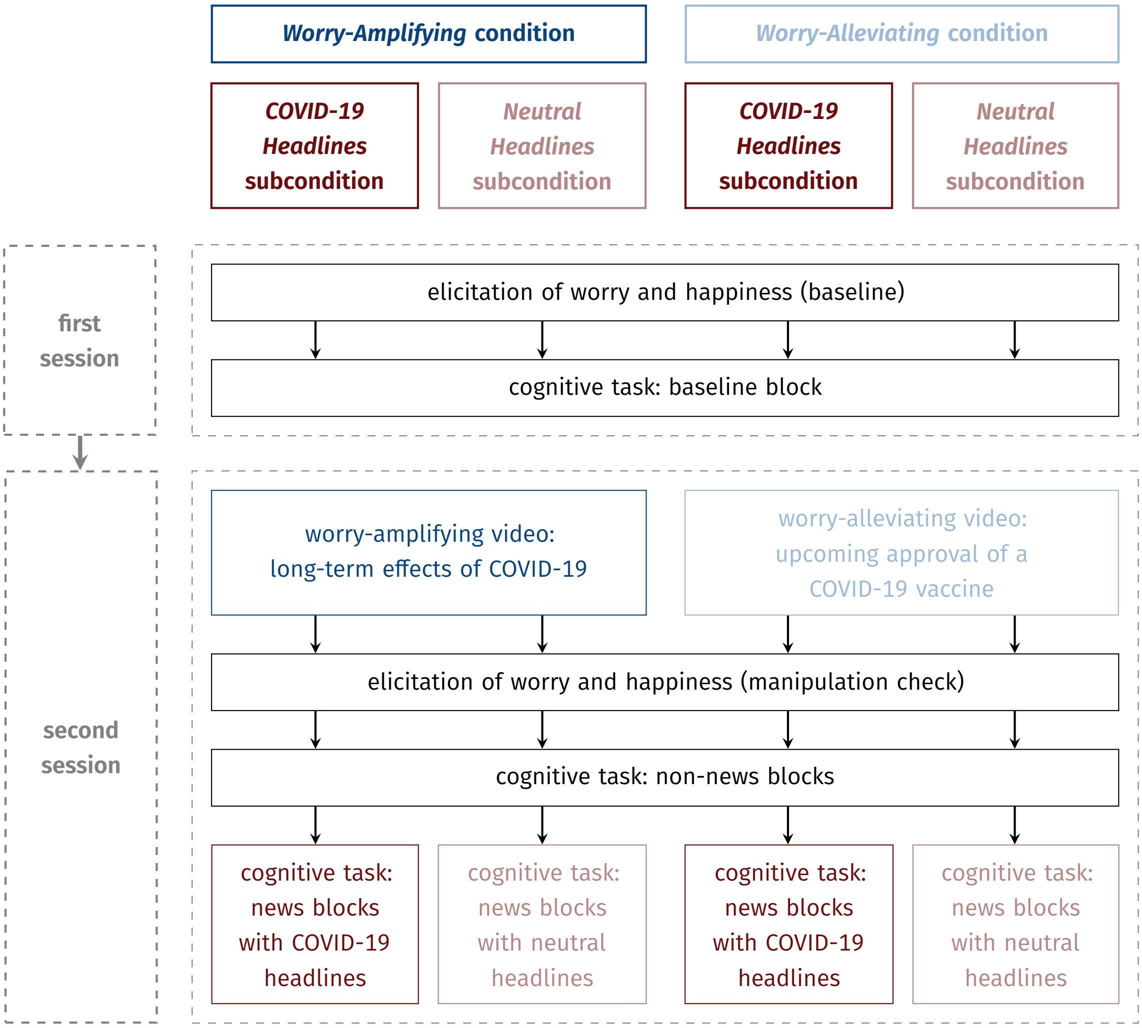

Figure 3.1. Chronology of the Main Components of the Experimental Design by Experimental Condition

Notes: Timeline of the most important components of the experimental design. For each combination of main and secondary experimental conditions, arrows indicate the sequence of experimental components. Sets of components that belong to the same session are framed by dashed gray rectangles. Within each session, experimental components that are identical in all conditions are drawn in black and span the whole width of the respective dashed rectangle. Components that differ between conditions *Worry-Amplifying* and *Worry-Alleviating*, constituting the main treatment manipulation, are highlighted in dark and light blue, respectively, whereas components that differ between the *COVID-19 Headlines* and *Neutral Headlines* subconditions are colored in dark and light red.

the main treatment manipulation and the cognitive task. In particular, (i) subjects are told at the beginning of the second session that they have to complete it without interruptions, (ii) instructions and comprehension questions for all parts of the second session are scheduled before the treatment manipulation, (iii) a pop-up window prompts subjects to continue with the experiment if they take longer than three minutes for the manipulation check between the news video clip and the cognitive task, and (iv) the first block of mental arithmetic problems starts automatically once subjects have completed the manipulation check.[12]

3.3 Preregistered Hypotheses and Analyses

My experiment is designed to facilitate tests of three main hypotheses relating to the effect of worrying about COVID-19 on labor productivity. All hypotheses are preregistered in a pre-analysis plan that was uploaded to the AEA RCT Registry before the start of the experiment. The preregistration documents are accessible in the registry under identifying number AEARCTR-0006558.[13]

As explained in the pre-analysis plan, all analyses are conducted using subject-level ordinary least squares (OLS) regressions with robust standard errors if not otherwise specified. Whenever possible, I control for the respective outcome measure at baseline to increase statistical power. For task outcomes, this is defined as the respective outcome in the second mental arithmetic block of the first session. I intentionally do not use data from the first block because this is the first time participants try out the task, so I expect additional noise in this data. Covariates are centered around the sample mean for continuous variables and around the sample median for non-binary discrete variables to facilitate interpretation of the constant. For most analyses, this implies estimation equations of the form

$$Y_i \ = \ \alpha + \beta \mathit{Worry\text{-}Amplifying}_i + \gamma' \boldsymbol{X}_i + \epsilon_i, \tag{3.1}$$

where Y_i is the outcome of interest for subject i, $\mathit{Worry\text{-}Amplifying}_i$ is an indicator equal to one if subject i is in the *Worry-Amplifying* condition, and $\boldsymbol{X}_i$ is a vector of control variables that contains the respective baseline outcome and may include additional covariates. The coefficient of interest is β.

3.3.1 Primary Hypotheses

The first hypothesis concerns the effect of the main treatment manipulation on subjects' level of worry about the health consequences of COVID-19 and constitutes a manipulation check. I expect that the main treatment manipulation shifts subjects' worry about COVID-19. This is a prerequisite for all subsequent analyses.

Hypothesis 3.1. *The level of worry about the health consequences of COVID-19 for oneself and one's loved ones is higher in condition Worry-Amplifying than in condition Worry-Alleviating.*

12. An analysis of participants' response times shows that these measures were successful: 91.8 percent of subjects spend less than 60 seconds longer than the video clip duration on the video page, and 98.5 percent of subjects spend less than three minutes on the manipulation check, i.e., less than 60 seconds per question.

13. The names of the experimental conditions were changed after running the experiment. The preregistration documents still contain the old condition names. The main experimental conditions were previously called *High concern* and *Low concern* instead of *Worry-Amplifying* and *Worry-Alleviating*, while the secondary conditions were called *High news salience* and *Low news salience* in place of *COVID-19 Headlines* and *Neutral Headlines*.

Since worry is not an isolated emotional response, it could be accompanied by other negative emotions. To assess the extent to which changes in productivity should be interpreted as the effect of worry in particular or as the effect of negative emotional responses to epidemics more generally, I also check whether the video affects feelings of happiness.

My main hypothesis concerns the effect of worrying on labor productivity. It is based on studies from cognitive psychology, which indicate that worrying can divert scarce attentional resources towards the perceived threat, thus reducing the available cognitive capacity for threat-unrelated processing tasks (Mathews, 1990; Eysenck et al., 2007; Robinson et al., 2013; Moran, 2016; Sari, Koster, and Derakshan, 2017). Intuitively, intrusive thoughts of worry can induce cognitive load in the same way as external distractions like a honking car or an overheard conversation. In line with this mental model, subjects exposed to the worry-evoking video clip should experience short-term reductions in their cognitive function that decrease their labor productivity for a given level of effort. To the extent that the emotional response to the video also includes a reduction in happiness, the effect of happiness on productivity documented in Oswald, Proto, and Sgroi (2015) suggests a complementary dynamic.[14]

Hypothesis 3.2. *Labor productivity, measured as the average number of correctly solved mental arithmetic problems across non-news blocks, is lower in condition Worry-Amplifying than in condition Worry-Alleviating.*

The intuition behind Hypothesis 3.2 also clarifies that there are two contingencies under which it might be untrue: (i) if worrying does not affect cognitive function in the current setting, e.g., because it is context-specific and subjects have already adapted to COVID-19 eight months into the pandemic, or (ii) if subjects compensate for the cognitive effects of worrying by putting in more mental effort, e.g., because they want to attain a goal number of correct answers that serves as a salient reference point or an income target (e.g., Heath, Larrick, and Wu, 1999; Fehr and Goette, 2007).[15]

Lastly, I hypothesize that the effect of worry on labor productivity is exacerbated by a continued salience of the pandemic in news headlines. The underlying intuition is that worry-induced distractions are temporary: even for individuals who are very scared of the disease,

14. I do not attempt to draw a conceptual distinction between the cognitive processes behind the productivity effects of changes in worry and happiness. In fact, Oswald, Proto, and Sgroi (2015) mention the possibility that increases in happiness free up cognitive resources otherwise captured by worrying on everyday problems as one potential mechanism behind their findings. Moreover, I do not take a stance on or investigate empirically to what extent these emotional responses and the resulting cognitive effects are productive in other ways, e.g., by increasing the chances of surviving the threat or providing some sort of catharsis for one's anxiety.

15. A review of the current state of research on the role of reference dependence for effort choices is provided by Goette (2021). Most relevant to my setting are studies that document increases in effort for individuals at risk of falling short of a personal goal. For instance, marathon runners increase their pace at the end of the race if they would otherwise just miss a round finishing time like the four-hour mark (Allen, Dechow, Pope, and Wu, 2017) and report reduced satisfaction if they fail to stay within their stated time goal (Markle, Wu, White, and Sackett, 2018). Moreover, bicycle taxi drivers in Kenya work more on days on which they report higher cash needs (Dupas, Robinson, and Saavedra, 2020), and Singaporean taxi drivers respond to unexpected booking cancellations or passenger no-shows by increasing their subsequent work productivity to offset the loss in earnings (Duong, Chu, and Yao, Forthcoming).

there will be phases during which worry fades and the mind can focus on the task. I expect that continued salience of the pandemic prevents this by keeping worry top-of-mind.[16] However, this hypothesis should be viewed as exploratory because my sample size is not large enough to draw a firm conclusion about it.

Hypothesis 3.3. *The negative effect of the Worry-Amplifying condition on labor productivity is stronger in subcondition COVID-19 Headlines than in subcondition Neutral Headlines.*

To test Hypothesis 3.3, I augment Equation 3.1 by two additional regressors and estimate a difference-in-differences equation. In particular, I regress productivity in news blocks on indicators for the *Worry-Amplifying* condition, the *COVID-19 Headlines* subcondition, and their interaction, controlling for baseline productivity. My hypothesis corresponds to a negative coefficient on the interaction term.

3.3.2 Other Preregistered Analyses

I also preregistered a number of additional analyses to gather supporting evidence and learn more about the determinants of the hypothesized relationships. I briefly summarize these analyses here.

(1) I test for an effect of the main treatment manipulation on subjects' beliefs about the danger of COVID-19. If the news video clips affected worry by shifting beliefs rather than conveying feelings, I would expect beliefs to differ across the two main experimental conditions.
(2) I assess the evolution of the main treatment effect over time within the four mental arithmetic blocks of the second session, focusing on participants in the *Neutral Headlines* subcondition.
(3) I test for an adverse effect of the main treatment manipulation on three measures of specific cognitive processes that determine productivity on the cognitive task: reaction time to the appearance of new mental arithmetic problems, calculation time per problem, and error rate.
(4) I investigate whether greater changes in worry or happiness in response to the main treatment manipulation are associated with greater reductions in productivity relative to baseline.
(5) I conduct heterogeneity analyses with respect to PSWQ score and self-reported exposure to media coverage about the pandemic. I expect a stronger effect for participants with a greater tendency to worry. I do not have a prior expectation regarding the heterogeneous

16. Mullainathan and Shafir (2013, chapter 2) report the results of a related experiment that illustrates how everyday cues can trigger scarcity-induced distractions. In their experiment, they compare the performance of dieters and non-dieters on a word search task with a neutral word (e.g., "cloud") after either having done a word search task with a tempting word (e.g., "cake") or a neutral word (e.g., "street") just before. They find that dieters on average take 30 percent longer to find the identical neutral word if the preceding task included a tempting word, whereas there is no effect for non-dieters. Interpreting this result, they argue that the tempting word brings a pre-existing longing for food top-of-mind, which then continues to capture subjects' attention and distracts them also on the next temptation-unrelated task.

effect for participants who follow the news about COVID-19. On the one hand, this could indicate less scope for the treatment to shift participants' level of worry, happiness and beliefs. On the other hand, those most concerned about the pandemic could be more likely to follow the news about it.

(6) I look for direct evidence that the main treatment manipulation has cognitive effects by testing for selective attention to pandemic-related news articles. In particular, I regress measures of news interest on indicators for the *Worry-Amplifying* condition, the *COVID-19 Headlines* subcondition, and their interaction. I control for the number of attempted tasks at baseline to proxy for the number of headlines participants see during the two news blocks of the task. I expect a positive coefficient on the interaction term, indicating that worried participants show a greater interest in news about the pandemic as compared to neutral news.

3.4 Results

Of the total number of 305 participants, 279 completed both sessions of the experiment.[17] Attrition is mostly driven by subjects whose devices did not satisfy the technical requirements for the task and subjects who did not show up to the second session. No subject dropped out after the main treatment manipulation, so differential attrition is not a concern. In line with the pre-analysis plan, five subjects were excluded from the sample because they did not pass the attention check after watching the video clip and six more were excluded because they did not solve a single mental arithmetic problems in at least one block, resulting in a final analysis sample of 268 subjects.

In the following, I first report baseline sample characteristics and show that the main treatment manipulation successfully shifted subjects' emotional responses to the COVID-19 pandemic. Then, I present results on the effect of the two treatment manipulations on labor productivity, in accordance with the hypotheses outlined in Section 3.3. Afterwards, I report suggestive evidence of changes in cognition and perceived productivity in response to the main treatment manipulation.

3.4.1 Baseline Sample Characteristics

Descriptive statistics for the analysis sample are reported in Table 3.A.1. One noteworthy observation is the relatively high average degree of worry at baseline, with a mean of 6.1 and a median of 7 on a scale from 0 to 10. This is reassuring because it implies that the vast majority of my student sample is sufficiently concerned about COVID-19 to respond to my experimental manipulations.

Table 3.A.2 presents baseline sample characteristics by main experimental condition. In addition to the mean and SD, I also report the standardized difference of means between the two conditions, defined as the difference in means normalized by the pooled sample SD. As

17. This includes two subjects who chose not to provide their bank details and therefore could not receive their payment despite finishing the experiment.

expected given randomization, the sample means of most baseline variables are not meaningfully different across conditions. However, the baseline level of worry in the *Worry-Amplifying* condition is almost 0.7 scale points lower than that in the *Worry-Alleviating* condition, corresponding to a standardized difference of means of 0.28 SD. Since this imbalance seems large enough to potentially distort some of my results, I additionally present estimates from specifications which control for baseline worry for all statistical analyses reported in the main text, even if this is not specified in the pre-analysis plan.

3.4.2 Manipulation Check: Emotional Responses to the Main Treatment Manipulation

Before turning to the effect of the experimental manipulations on labor productivity, I first check whether the main treatment manipulation succeeded in changing participants' emotional response to COVID-19, in line with Hypothesis 3.1. The coefficient estimates of this manipulation check are displayed in Table 3.1.

Table 3.1. Manipulation Check: Effect of the Main Treatment Manipulation on Worry and Happiness

Dependent variable:	Change in worry	Worry		Change in happiness	Happiness	
	(1)	(2)	(3)	(4)	(5)	(6)
Worry-Amplifying condition	0.819***	0.152	0.731***	−1.473***	−1.405***	−1.446***
	(0.142)	(0.286)	(0.135)	(0.224)	(0.247)	(0.201)
Worry at baseline (cent.)			0.868***			
			(0.032)			
Happiness at baseline (cent.)						0.592***
						(0.063)
Constant	−0.083	6.348***	6.842***	0.341**	6.530***	7.010***
	(0.079)	(0.198)	(0.079)	(0.154)	(0.164)	(0.119)
Observations	268	268	268	268	268	268
R^2 (adjusted)	0.107	−0.003	0.765	0.137	0.105	0.403
Mean (dependent variable)	0.332	6.425	6.425	−0.407	5.817	5.817

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

In column (1), the dependent variable is the change in worry from baseline to after the news video clip. On average, there is a significant increase in worry in the *Worry-Amplifying* condition of about 0.8 scale points, while the decrease in worry in the *Worry-Alleviating* condition is small and insignificant. This suggests that the main treatment manipulation successfully shifted subjects' worry about COVID-19, and this was driven by the worry-inducing video clip.

In columns (2) and (3), I test whether the differential shift in worry after the treatment manipulation translates into level differences across the two conditions, once without and once with adjustment for worry at baseline. The raw difference in average posttreatment worry across the two conditions is small and insignificant. This is attributable to the baseline imbalance: the increase in worry in the *Worry-Amplifying* condition closes the initial gap between the two conditions rather than opening up a new one. The preregistered specification in column (3) controls for baseline worry and therefore accounts for the baseline imbalance. In this specification, the worry gap between the two conditions increases to about 0.7 scale points, which is highly significant and captures about 80 percent of the treatment-induced change in worry from the first to the second session estimated in column (1).

All in all, these estimates provide support for Hypothesis 3.1. As long as the initial imbalance is accounted for by controlling for worry at baseline, participants in the *Worry-Amplifying* condition exhibit an exogenously higher level of worry compared to those in the *Worry-Alleviating* condition, driven by a large increase in worry in the second session. In Appendix 3.B, I investigate whether the shift in worry is driven by new factual information contained in the video clips or by the communicated feelings and emotions. For this purpose, I estimate the effect of the main treatment manipulation on elicited beliefs about the virus. The identified pattern of belief shifts suggests that the effect is at least partly driven by conveyed emotions rather than information provision.

The remaining columns contain analogous coefficient estimates for happiness. This reveals whether the increase in worry is accompanied by a reduction in happiness, which could trigger complementary—though presumably conceptually similar—cognitive processes through which epidemic-induced emotions impair productivity. The estimates show a gap in happiness of about 1.4 scale points between conditions *Worry-Amplifying* and *Worry-Alleviating* after the main treatment manipulation, which is mainly driven by a large decrease in happiness following the worry-evoking video. This implies that the results in the subsequent subsections should be interpreted as the joint effect of a multi-faceted negative emotional response to the pandemic.

3.4.3 Effect of the Main Treatment Manipulation on Labor Productivity

In Table 3.2, I present coefficient estimates from tests of Hypothesis 3.2, which predicts that labor productivity decreases after watching the worry-inducing news clip about COVID-19. The dependent variable is the average number of correct answers across the two non-news blocks of mental arithmetic problems, before the start of the secondary treatment manipulation. I report three specifications which differ in the included covariates. The regression results offer no evidence that the main treatment manipulation affects productivity. The point estimate of my preferred specification in column (3), which controls for both worry and productivity at baseline, indicates that exposure to the fear-evoking video causes a drop in the number of correct answers of 0.137. Relative to the mean of the *Worry-Alleviating* condition, this is a decrease of about 1.1 percent. The estimated treatment effect is far from significant in both statistical and economic terms as well as reasonably precise. In particular, I can comfortably rule out a 0.2 SD effect at the five percent significance level.[18]

18. The partitioning of the cognitive task into blocks with independent sequences of mental arithmetic problems within each block also permits an investigation the evolution of the treatment effect over time. Figure 3.A.1 in Appendix 3.A plots separate coefficient estimates of the main treatment effect for each block of mental arithmetic

Table 3.2. Effect of the Main Treatment Manipulation on Labor Productivity in Non-news Blocks

	Dependent variable: Correct answers per non-news block		
	(1)	(2)	(3)
Worry-Amplifying condition	0.069	−0.069	−0.137
	(0.478)	(0.280)	(0.286)
Correct answers at baseline (cent.)		0.799***	0.791***
		(0.033)	(0.033)
Worry at baseline (cent.)			−0.105
			(0.065)
Constant	12.023***	11.326***	11.274***
	(0.342)	(0.193)	(0.201)
Observations	268	268	268
R^2 (adjusted)	−0.004	0.657	0.660
Mean (dependent variable)	12.058	12.058	12.058

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.C.1 in Appendix 3.C reports estimates from alternative specifications that also include demographic control variables or account for baseline emotions nonlinearly. The estimated treatment effect is slightly greater in absolute size in some of these specifications, ranging up to -0.268. However, it always remains statistically insignificant and economically small. In light of the documented baseline imbalance in worry, it is especially reassuring that the inclusion of a set of dummies for the possible values of baseline worry does not substantially change the results.

The regression results in Table 3.2 show that the main treatment manipulation does not have an effect on productivity on average, but they might not reveal more general distributional effects. Therefore, Figure 3.2 additionally plots the empirical cumulative distribution of labor productivity by main experimental condition, once before and once after adjusting for covariates. The resulting distributions look similar in both panels of the figure. To assess formally whether there are significant differences in the shape of the outcome distribution across conditions, I conduct Kolmogorov-Smirnov tests. Since the standard version of the test suffers from low statistical power in the presence of ties in the data, I use a randomization test with 10 000 repetitions to compute the p-values manually (Neuhäuser, Welz, and Ruxton, 2017). The underlying idea is that under the null hypothesis that the treatment effect is zero for all subjects, the value of the Kolmogorov-Smirnov test statistic should be unaffected by random permutations of participants' treatment status. To obtain a p-value, I compare the observed value of the Kolmogorov-Smirnov test statistic to the randomization distribution of test statistic values that

problems of the second session. Although the coefficient for the first block is slightly greater in absolute size, none of the estimates is significantly different from zero.

emerges from reshuffling treatment status and recalculating the test statistic a large number of times. An observed value of the test statistic in the tails of the randomization distribution is evidence against the null hypothesis. Rosenbaum (2002) shows that randomization inference can also be applied to covariate-adjusted outcomes. The resulting p-values of the Kolmogorov-Smirnov randomization tests are 0.809 and 0.712 for the raw and covariate-adjusted outcome distributions, respectively. Thus, I cannot reject the null hypothesis that the distribution of labor productivity is the same in both experimental conditions.

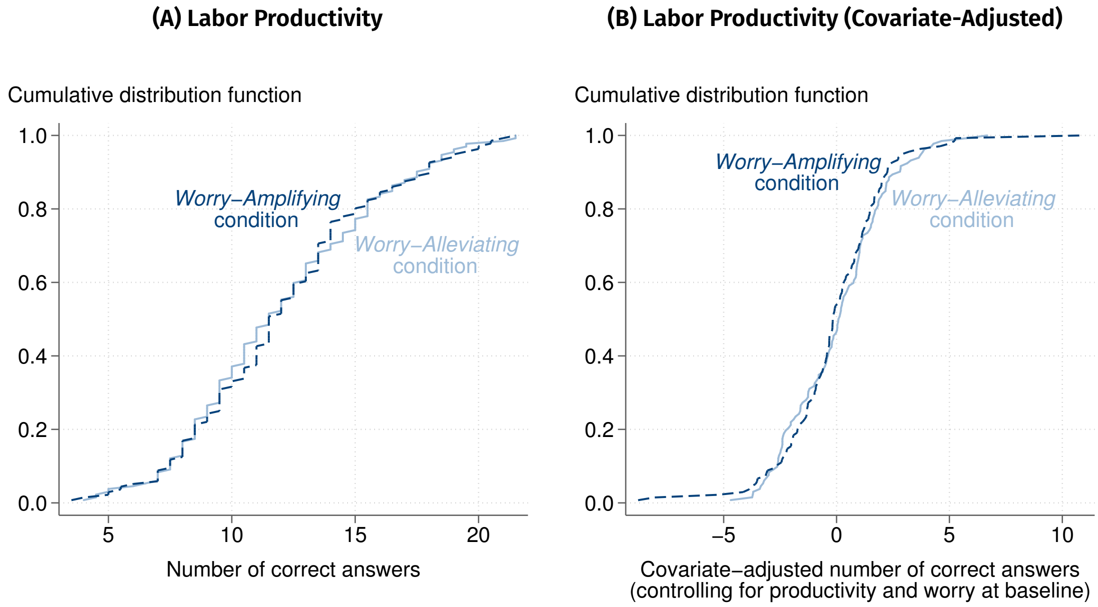

Figure 3.2. Labor Productivity in Non-news Blocks by Experimental Condition

Notes: The cumulative distribution of labor productivity in non-news blocks, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of the average number of correct answers across the two non-news blocks of mental arithmetic problems. Panel (B) shows the distribution of the residuals from a regression of the average number of correct answers on the baseline number of correct answers and the baseline level of worry.

I also test whether the main treatment manipulation affects measures of attention and working memory, two specific cognitive processes that determine productivity on the cognitive task. The respective coefficient estimates are reported in Table 3.A.3 and Table 3.A.4 in Appendix 3.A.[19] In line with the above results for the primary productivity measure, the estimates provide no evidence that watching the fear-evoking video has an impact on subjects' reaction time to the appearance of new mental arithmetic problems, their average calculation time per problem, or their error rate.

Next, I investigate the relationship between within-subject changes in productivity and changes in emotional state. If productivity is impaired by negative emotional reactions to COVID-19, then subjects with stronger emotional responses to the main treatment manipula-

19. Cumulative distribution plots for the individual cognitive processes can be found in Figure 3.A.2 and Figure 3.A.3.

tion should experience greater decreases in productivity in non-news blocks relative to baseline. Table 3.3 contains estimates from regressions of changes in productivity on changes in worry and happiness.[20] All point estimates are close to zero and far from statistically significant, indicating that there is no association between treatment-induced changes in emotional state and changes in productivity.

Table 3.3. Changes in Worry and Happiness and Changes in Labor Productivity

	Dependent variable: Change in productivity		
	(1)	(2)	(3)
Change in worry	0.050		0.062
	(0.105)		(0.107)
Change in happiness		0.025	0.034
		(0.072)	(0.074)
Constant	1.082***	1.109***	1.092***
	(0.153)	(0.153)	(0.157)
Observations	268	268	268
R^2 (adjusted)	−0.003	−0.003	−0.006
Mean (dependent variable)	1.099	1.099	1.099

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Change in productivity* is the difference in the average number of correct answers per block between non-news blocks and baseline. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Even if the treatment has no meaningful average effect in the full sample, it could still influence the labor productivity of especially susceptible subgroups. Two such subgroups were specified in the pre-analysis plan. First, I expect subjects with a higher tendency to worry, as revealed by a higher PSWQ score, to respond more strongly to the treatment manipulation. Second, one might expect differential effects by subjects' level of attention to media coverage about COVID-19. On the one hand, this could be a sign of strong pre-existing concerns about COVID-19. On the other hand, it could indicate less scope for the treatment manipulation to shift participants' emotions because their opinions about the pandemic are less malleable.

Both candidate characteristics are related to baseline differences in emotional state. Subjects with higher PSWQ scores are more worried and less happy at baseline, with differences in means of 0.809 ($p = 0.005$, two-sided test with robust standard errors) and 1.433 ($p < 0.001$) scale points between subjects with above-median and below-median scores, respectively. Subjects who report following the news about COVID-19 are more worried, but equally happy, at baseline compared to their less interested peers. Here, the difference in the mean level of worry across the two groups is 0.759 scale points ($p = 0.022$).

20. I also compute Spearman's rank correlation coefficients, presented in Table 3.A.5. With respect to the relationship between changes in productivity and changes in emotions, the estimated correlations support the conclusions of the regression analysis. An additional insight of the correlation matrix is that changes in worry are negatively correlated with changes in happiness with a correlation coefficient of -0.165 ($p = 0.007$), suggesting that subjects' negative emotional response to the treatment manipulation in these two dimensions is closely linked.

To investigate whether these baseline differences in emotional state translate into differential treatment effects, I conduct heterogeneity analyses. Figure 3.3 plots separate estimates of the main treatment effect for different partitions of the analysis sample along the two specified characteristics.[21]

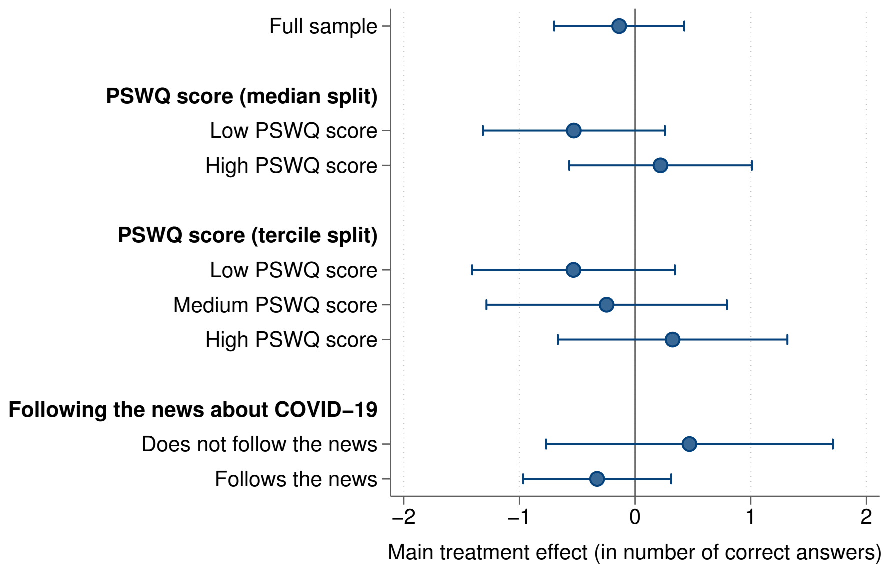

Figure 3.3. Heterogeneity in the Main Treatment Effect on Labor Productivity

Notes: Plot of coefficient estimates of the main treatment effect on labor productivity for the full sample and three different partitions of the full sample into subsamples: a median split by PSWQ score, a tercile split by PSWQ score, and a split by whether or not subjects report following the news about COVID-19. The estimates are from regressions of the subject-level number of correct answers in non-news blocks on indicators for each of that partition's subsamples and their interactions with an indicator for the *Worry-Amplifying* condition, controlling for the number of correct answers and the level of worry at baseline. Error bars indicate 95 percent confidence intervals, constructed using heteroscedasticity-robust standard errors.

Qualitatively, the point estimates suggest stronger treatment effects for subjects with lower PSWQ scores, as if those with a greater general tendency to worry were more adept at warding off worry-induced distractions, and for subjects who follow the news about COVID-19. However, the estimates are always above the lower confidence limit of the full sample treatment effect and never significantly different from zero. Thus, there is no evidence that worrying about COVID-19 impairs productivity even among assumably more susceptible subpopulations.[22]

21. I also show results from a tercile split by PSWQ score to alleviate concerns that heterogeneous effects might be non-monotonic.

22. I also check for heterogeneous emotional responses to the treatment manipulation. The respective estimates, plotted in Figure 3.A.4 in the appendix, show no clear patterns of stronger emotional responses for more susceptible subgroups of the sample. The treatment manipulation induces significant shifts in both worry and happiness in all subsamples. Subjects with a lower PSWQ score seem to show relatively weaker emotional responses, but only the differential treatment effect in happiness of 0.952 scale points relative to the medium-PSWQ-score subsample is significant at the ten percent level ($p = 0.057$ in an F-test).

Taking stock, I find no evidence that the main treatment manipulation reduces labor productivity across a variety of conducted analyses even though it induces a strong negative emotional response. In a final analysis regarding Hypothesis 3.2, I test whether unintended side effects of the intervention which are unrelated to worrying about COVID-19 could mask an effect of worrying on productivity. A prime candidate is that the worry-evoking video motivates subjects to work harder by shifting their beliefs about the financial consequences of the pandemic. In particular, learning about the severity of long-term health effects could induce pessimism about the duration of containment measures, whereas faster vaccination campaigns signal a faster economic recovery. Thus, the two media reports might shift subjects' beliefs about their own future economic prospects in opposite directions, ultimately resulting in a greater perceived marginal value of each Euro earned in the experiment for subjects exposed to the worry-inducing video. As a consequence, subjects in the *Worry-Amplifying* condition might work harder for reasons unrelated to their negative emotional response, thereby counterbalancing an adverse effect of worrying on productivity. Kaur, Mullainathan, et al. (2021) find evidence of such an offsetting motivational effect of a priming intervention designed to increase financial worries in a field experiment with Indian workers under financial strain.

My collected data allows me to assess the plausibility of this concern. Table 3.4 contains estimates from regressions on an indicator for the *Worry-Amplifying* condition that use subjects' reported level of worry about the consequences of the pandemic on their personal financial situation or their job prospects as the dependent variable. Moreover, I also test whether the main treatment manipulation affects the goal number of correct answers per mental arithmetic block that subjects state just after watching their condition's video clip. None of the reported specifications provides any indication that the main treatment manipulation shifts subjects' beliefs about their financial prospects or motivates them to set higher goals for themselves. The estimated treatment effects for participants' financial concerns are substantially smaller than the respective estimates for health-related worry and happiness in Table 3.1, and those for their productivity goals have the wrong sign. Moreover, none of the estimates is significantly different from zero.

Since the lack of a treatment effect is not explained by countervailing side effects of the treatment manipulation, the natural conclusion is that emotional responses to the COVID-19 pandemic do not impair labor productivity in my experiment. The intuition behind Hypothesis 3.2 suggests two potential explanations for this null result. First, worrying might not impede cognitive ability in my setting. Eight months into the pandemic, subjects in my sample might have already adapted to the situation and the accompanying negative emotions, such that emotional responses provoked by media reports are not extreme enough to induce a capture of cognitive resources. Second, worry might impede cognitive ability, but subjects compensate for worry-induced cognitive effects by increasing their mental effort.[23] This would be in line with the notion of income targeting: subjects might have a goal number of correct answers in

23. There is a subtle difference between this potential explanation for the observed null result and the possibility that an adverse effect of worrying is masked by an offsetting motivational side effect of the treatment manipulation, as explored in Table 3.4. In particular, only the latter would predict a reduction in labor productivity in situations that arouse worry about COVID-19 without shifting beliefs about financial prospects.

Table 3.4. Effect of the Main Treatment Manipulation on Selected Outcomes

Dependent variable:	Worry about financial situation		Worry about job prospects		Goal for correct answers per block	
	(1)	(2)	(3)	(4)	(5)	(6)
Worry-Amplifying condition	−0.008	0.166	−0.093	0.083	−0.280	−0.340
	(0.330)	(0.325)	(0.356)	(0.352)	(0.575)	(0.451)
Worry at baseline (cent.)		0.261***		0.264***		0.133
		(0.071)		(0.077)		(0.135)
Correct answers at baseline (cent.)						0.860***
						(0.062)
Constant	3.538***	3.686***	3.917***	4.066***	11.530***	10.856***
	(0.245)	(0.244)	(0.248)	(0.250)	(0.366)	(0.233)
Observations	268	268	268	268	268	268
R^2 (adjusted)	−0.004	0.045	−0.003	0.038	−0.003	0.511
Mean (dependent variable)	3.534	3.534	3.869	3.869	11.388	11.388

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

mind that acts as their reference point, and they react to unexpected decreases in perceived productivity by increasing their mental effort to avoid falling short of their goal (e.g., Allen et al., 2017; Duong, Chu, and Yao, Forthcoming). This behavior might even be promoted by the subtle experimental design feature of asking subjects for their goal number of tasks before the second session's working period.

3.4.4 News Salience and Labor Productivity

Finally, I turn to Hypothesis 3.3, which predicts that the effect of worry on labor productivity is compounded by a continued salience of the pandemic because this keeps worry top-of-mind. In light of the previous finding that worrying does not affect productivity in the first place, there is no conceptual reason to believe that keeping it top-of-mind should matter. For the sake of completeness, I nonetheless report the corresponding estimates in Table 3.A.6 and Table 3.A.7 in the appendix.

To test the hypothesis, I analyze labor productivity in news blocks, during which news headlines are displayed to subjects in between mental arithmetic problems. The headlines are either about the COVID-19 pandemic or neutral topics, depending on the cross-randomized secondary experimental condition. I report estimates of the effect of the *COVID-19 Headlines* subcondition for each of the two subsamples that arise from the main treatment manipulation as well as the interaction between indicators for conditions *COVID-19 Headlines* and *Worry-Amplifying* for the full sample. The interacted model additionally accounts for productivity effects of the displayed news headlines that are unrelated to worrying about COVID-19. For instance, the neutral headlines could be more interesting and thus more distracting than the pandemic-related headlines even if worrying plays no role. All models are estimated once with and once without additionally controlling for worry at baseline.

In line with expectations given the previous results, the estimated coefficients of interest are close to zero and statistically insignificant in all models.

3.4.5 Evidence for Worry-Induced Attentional Capture

The intuition behind the hypothesized effect of worrying on labor productivity is that scarce cognitive resources are redirected towards the object of concern, thus reducing their availability for other tasks (e.g., Eysenck et al., 2007; Mullainathan and Shafir, 2013; Lichand and Mani, 2020). Therefore, one potential explanation for the observed lack of a treatment effect in my setting is that the emotional responses provoked by the news video clips in my experiment simply do not affect subjects' cognitive processes. To get a first indication of the empirical plausibility of this explanation, I look for direct evidence of changes in cognition in response to the main treatment manipulation.

Remember that participants are exposed to headlines from current news articles in between mental arithmetic problems during the two news blocks of the second session. Depending on their secondary experimental condition, the displayed headlines are either about a neutral topic or related to the COVID-19 pandemic. To proxy for attentional capture, I use participants' interest in the news headlines that they are exposed to, as measured by whether they click on a button to save them for later reading. In particular, I examine whether the relative propensity to save news articles about COVID-19 differs between the *Worry-Amplifying* and *Worry-Alleviating* condition.

Corresponding regression estimates are reported in columns (1) and (2) of Table 3.5. I account for differences in the number of headlines subjects see by adding the attempted number of arithmetic problems at baseline as an additional covariate. The small and insignificant point estimate on the coefficient for the *COVID-19 Headlines* subcondition indicates that the article topic barely affects the number of saved articles in the *Worry-Alleviating* condition. In comparison, the worry-inducing video clip reduces the number of neutral articles bookmarked by statistically significant 1.07 articles, which is a 50 percent reduction relative to the average number of saved neutral articles in the *Worry-Alleviating* condition of 2.30.[24] At the same time, the worry-inducing video clip does not meaningfully reduce the number of pandemic-related articles that subjects save, as implied by a slightly smaller interaction effect of opposite sign. This interaction effect, however, is not statistically significant.

As revealed by the small value of the constant, the general propensity of experimental participants to save articles is low. As a consequence, few subjects who save many articles could have a large impact on the estimates in columns (1) and (2). To assess whether the conclusions

24. I did not specify a hypothesis for this coefficient in the pre-analysis plan because the main treatment manipulation could also have a topic-independent effect on the number of saved articles, e.g., by reducing the number of mental arithmetic problems subjects attempt. However, in the absence of a treatment effect on the number of attempted problems—which I verify in unreported regressions—, the large negative coefficient already provides a first indication that worrying reduces attention to pandemic-unrelated aspects of information.

Table 3.5. Effect of the Main Treatment Manipulation on News Interest

Dependent variable:	Number of saved articles		At least one saved article	
	(1)	(2)	(3)	(4)
Worry-Amplifying condition	−1.115**	−1.070**	−0.295***	−0.286***
	(0.469)	(0.482)	(0.083)	(0.084)
COVID-19 Headlines subcondition	−0.101	−0.100	−0.086	−0.085
	(0.529)	(0.531)	(0.088)	(0.088)
Worry-Amplifying condition × *COVID-19 Headlines* subcondition	0.931	0.933	0.324***	0.324***
	(0.700)	(0.701)	(0.122)	(0.122)
Attempted answers at baseline (cent.)	0.002	0.010	−0.006	−0.004
	(0.043)	(0.043)	(0.009)	(0.009)
Worry at baseline (cent.)		0.073		0.015
		(0.076)		(0.013)
Constant	2.255***	2.297***	0.594***	0.603***
	(0.378)	(0.379)	(0.061)	(0.063)
Observations	268	268	268	268
R^2 (adjusted)	0.010	0.010	0.041	0.043
Mean (dependent variable)	1.866	1.866	0.481	0.481

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. The dependent variables are measures of subjects' interest in task-irrelevant news articles whose headlines are displayed to them in between mental arithmetic problems during news blocks: the total number of articles saved for later reading in columns (1) and (2) and an indicator for saving at least one article in columns (3) and (4). *Worry-Amplifying condition* and *COVID-19 Headlines subcondition* are indicators equal to one for subjects in the experimental conditions of the respective name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

are robust, I also use an indicator for saving at least one article as an alternative outcome measure. The resulting estimates in columns (3) and (4) are qualitatively in line with those for the count outcome and exhibit similarly meaningful effect sizes. In these specifications, the interaction effect is significantly different from zero as well.

Since the effect of the worry-inducing video on news interest depends on the article topic, it cannot be explained by a general effect of worrying on the ability to withstand distractions. Similarly, differences in the demand for information about COVID-19 induced by the main treatment manipulation cannot explain the observed lower interest in neutral news articles because these articles do not have any connection to the video clips. Therefore, the estimates suggest that exposure to the worry-inducing video indeed increases subjects' relative focus on pandemic-related news, in line with attentional capture. In particular, they document a substantial decrease in demand for news about topics unrelated to COVID-19. With respect to the interpretation of the null result for Hypothesis 3.2, this piece of evidence is of course only suggestive because I cannot assess whether the relative shift in subjects' attention towards pandemic-related news really translates into reduced focus on the cognitive task.

The estimated effects of the main treatment manipulation on news interest are also interesting in their own right. In text analyses of the media coverage of the COVID-19 pandemic in the US, Sacerdote, Sehgal, and Cook (2020) document that articles by major US news outlets are overwhelmingly negative in tone and the choice of covered topics, irrespective of real-world epidemiological developments. My estimates offer a potential explanation for this observation. They indicate that for media outlets which compete for the attention of worried readers, producing pessimistic articles about the pandemic might be an optimal strategy.[25]

3.4.6 Effect of the Main Treatment Manipulation on Self-Perceived Productivity

In Section 3.4.3 and Section 3.4.5, I show that worrying does not reduce labor productivity in the context of my experiment even though it shifts relative attention to news articles about COVID-19. In this subsection, I present suggestive evidence on one potential explanation for these findings: since solving mental arithmetic problems also has an effort component, subjects might be able to compensate for worry-induced reductions in task-available cognitive function by putting in more mental effort.[26]

One implication of an increase in required mental concentration is that the cognitive task should feel more difficult for subjects who are exposed to the worry-inducing treatment manipulation. Table 3.6 contains estimates of the effect of the main treatment manipulation on three different survey outcomes that are related to subjects' perceptions of their performance on the cognitive task.

In columns (1) and (2), the outcome is a summary measure of subjects' self-assessed productivity in non-news blocks and news blocks relative to the baseline session.[27] The dependent variable is coded as -1, 0, and 1 if subjects guess that they answered more, approximately the same, or fewer problems correctly in the second session, respectively.[28] The estimated coefficients indicate that perceived productivity is marginally significantly lower in the *Worry-Amplifying* condition for non-news blocks, while the estimated treatment effect is close to zero for news blocks.

In column (3), I compare subjects' reported level of satisfaction with their task performance in the second session across conditions. Taken at face value, the negative treatment effect estimate of one quarter of a scale point suggests that subjects who are exposed to the

25. Since my experimental design only varies the topic of displayed headlines, I cannot say whether positively framed headlines about COVID-19 would have generated similar increases in relative news interest. However, the selected pandemic-related headlines are arguably more pessimistic on average than those displayed in the *Neutral Headlines* subcondition.

26. The analyses in this subsection are not preregistered.

27. To keep the question simple, subjects are asked to consider their productivity relative to their performance in the whole first session rather than only in the second block of the first session, which is used as the baseline measure of task performance throughout the chapter. As a result, the definition of the term "baseline session" used in this subsection differs from the definition of baseline productivity used in other analyses.

28. Distributions of the individual responses by experimental condition are plotted in Figure 3.A.5. In general, participants have some insight into their productivity. The Spearman correlation coefficients between my measure of perceived productivity and the actual change in the number of correct answers relative to the first session is 0.386 ($p < 0.001$) for non-news blocks and 0.333 ($p < 0.001$) for news blocks.

Table 3.6. Effect of the Main Treatment Manipulation on Measures of Self-Perceived Productivity

Dependent variable:	Perceived productivity in non-news blocks	Perceived productivity in news blocks	Satisfaction with own performance	Concentration on the task impaired by video	
	(1)	(2)	(3)	(4)	(5)
Worry-Amplifying condition	−0.172*	0.023	−0.265	0.128***	0.157***
	(0.095)	(0.102)	(0.265)	(0.041)	(0.049)
Worry at baseline (cent.)	−0.018	−0.015	−0.042	0.029***	0.011
	(0.020)	(0.022)	(0.061)	(0.008)	(0.011)
Worry-Amplifying condition × Worry at baseline (cent.)					0.033**
					(0.017)
Constant	0.187***	−0.138*	6.794***	0.085***	0.074***
	(0.066)	(0.074)	(0.191)	(0.024)	(0.025)
Observations	268	268	268	268	268
R^2 (adjusted)	0.007	−0.005	−0.002	0.063	0.074
Mean (dependent variable)	0.116	−0.112	6.698	0.123	0.123

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. Perceived productivity is coded as −1, 0, and 1 if subjects guess that they answered more, approximately the same, or fewer mental arithmetic problems correctly on average in the named second-session blocks relative to the baseline-session blocks, respectively. *Concentration on the task impaired by video* is an indicator equal to one for subjects who report being distracted by thoughts about the content of their condition's video. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

worry-inducing video clip are slightly less satisfied with their performance on average, but the coefficient is not significantly different from zero.

Finally, I also ask subjects directly whether their concentration on the mental arithmetic problems was impaired by thoughts about the content of the video they had seen.[29] The coefficient estimates on subjects' answers to this question are reported in column (4). At the median value of baseline worry, the worry-inducing video clip increases the propensity to report being distracted by the content of the video by statistically significant 12.8 percentage points, which is an increase of 150 percent relative to the average of the *Worry-Alleviating* condition. Interestingly, the propensity to feel distracted by thoughts about the content of the video also increases in subjects' baseline level of worry about COVID-19, consistent with the notion that worrying makes it harder to concentrate on the cognitive task. This interpretation is further corroborated in column (5), which includes the interaction term between the treatment dummy and baseline worry. The likelihood of self-reported reductions in focus on the task caused by

29. Since this is a blunt question, it might be susceptible to experimenter demand effects (e.g., Zizzo, 2010). However, note that it is not obvious why experimenter demand should be stronger in the *Worry-Amplifying* condition. If subjects mistakenly believed that the main purpose of the experiment was to compare their task outcomes between the first and second session and wanted to please the experimenter to the same extent in both conditions, the estimated between-subjects treatment effect would be unaffected.

the worry-inducing video clip increases by about 3.3 percentage points for each scale point of baseline worry.

Taken together, the estimates point towards a decrease in subjects' perceived productivity caused by the worry-inducing video clip, in line with the idea that subjects feel less productive because they have to exert more effort to accomplish the cognitive task. However, the differences between the two conditions are only statistically significant at the five percent level for one of the four outcomes and thus far from conclusive.

3.5 Discussion

The preceding section documents that an exogenous multi-faceted negative emotional response to COVID-19, comprising an increase in worrying and a reduction in happiness, does not reduce labor productivity on a mental arithmetic task in a student sample. Being exposed to news headlines about the pandemic while working also does not have an effect on labor productivity. At the same time, there is suggestive evidence that worrying leads to a shift of attention towards pandemic-related news and makes subjects report a reduced level of focus on the task.

In this section, I discuss unaddressed concerns to the internal validity of these results and try to reconcile my findings with previous research.

3.5.1 Are the Results Affected by Cheating on the Cognitive Task?

While unavoidable during a pandemic, online experiments come with the disadvantage that they afford the experimenter less control about the behavior of participants during the experiment. In the current study, this raises a potential concern regarding the internal validity of my findings: subjects might cheat on the mental arithmetic problems, and cheating might crowd out treatment effects on productivity. My experimental design prevents the most obvious form of cheating—using a standard calculator—by keeping subjects' hands busy. However, inhibiting cheating completely on an online task that requires cognitive abilities is difficult. For instance, I cannot rule out that subjects circumvent my measures by getting help from a second household member or using a voice-controlled calculator. This possibility is also mentioned by one subject in the questionnaire. However, there are several patterns in the data that are inconsistent with widespread cheating.

First, there is a sizable number of incorrect answers. On average, subjects answer 13.5 percent of their attempted mental arithmetic problems incorrectly in the second session, and only 7 out of 268 subjects do not make a single mistake. If a majority of subjects used a voice-controlled calculator, I would expect fewer errors.

Second, there is evidence for a trade-off between improvements in calculation time and error rate. Naturally, subjects' performance on the task improves on average between the first and the second session as they gain practice. These improvements can be driven by decreases in the error rate or by decreases in calculation time, i.e., the number of seconds subjects spend thinking on each problem. However, one would normally expect large improvements in one dimension to come at the expense of the other dimension, reflecting the inherent trade-off between speed and quality. In contrast, subjects who come up with a way to cheat can probably achieve large improvements in both dimensions simultaneously. In support of rule compliance, the Pearson's correlation coefficient between the percentage change in success rate and the percentage change in average calculation time from baseline to the second session is 0.278 ($p < 0.001$).[30]

Third, subjects performance on individual problems varies with problem difficulty. Two of the first ten mental arithmetic problems in each block were deliberately selected to be easier than the others. Easy problems are characterized by fewer carries and borrows from the unit digit to the tens or hundreds digit or by a subtrahend that cancels out with one of the summands, thus reducing the number of mental operations required to arrive at the solution. An example is the problem $21 + 71 - 21 + 14$. For participants who cheat by using a voice-controlled calculator, the distinction between easy and hard problems is inconsequential because it does not meaningfully affect the speed of the calculator, but subjects who solve the problems by means of mental arithmetic should do better on easy problems. To test this, I compute the subject-level average calculation time and success rate in the second session separately for easy and more difficult problems. On average, subjects have a 29.8 percent lower average calculation time and a 7.3 percent higher success rate on easy problems (both significantly different from zero with $p < 0.001$ in a two-sided test). 171 subjects do strictly better on easy tasks in both dimensions, compared to only 11 subjects who do not improve in at least one.

To further alleviate concerns that cheating could affect my conclusions, I conduct a robustness check that excludes potential cheaters. In particular, I drop subjects who either improve by more than 20 percent in both calculation time and success rate from baseline to the second session or who do not do better on easy tasks in at least one dimension. These criteria also capture the one subject who brings up the possibility of using a voice-controlled calculator in the questionnaire. The estimated coefficients for the three main hypotheses, presented in Table 3.C.2 in Appendix 3.C, are very similar to those for the full sample.

Taken together, even though isolated cases of subject misbehavior are likely, there is evidence that the majority of participants adheres to the rules of the experiment, and the estimates are robust to the exclusion of suspected cheaters. Therefore, it seems unlikely that my conclusions are affected by the possibility of cheating on the cognitive task.

3.5.2 Differences to Findings from Previous Research

My findings relate to work in economics on the role of emotions as determinants of cognitive function and labor productivity. Further discussion seems merited with respect to two distinct lines of research: studies on the cognitive effects of worrying that originate from work on the psychological consequences of poverty, and studies on the effect of happiness on productivity from labor economics.

30. I use the success rate rather than the error rate here to avoid problems with division by zero for subjects without any errors at baseline.

Cognitive effects of worrying

The idea that worrying may impede cognitive function in general and labor productivity in particular is the focus of an emerging line of research in behavioral development economics on the psychology of poverty. It hypothesizes that poverty itself can impair decision-making and worker productivity among the poor because the perception of scarcity captures cognitive resources, thereby creating a vicious cycle of poverty (Shah, Mullainathan, and Shafir, 2012; Mullainathan and Shafir, 2013). Empirical evidence on the effect of financial strain on cognitive function has been mixed. While results from early priming experiments and analyses of natural variation in income before and after payday are generally supportive of the theory (e.g., Mani et al., 2013), they do not always replicate in more recent investigations in similar settings (e.g., Carvalho, Meier, and Wang, 2016). In an attempt to reconcile the existing evidence, Lichand and Mani (2020) show that income uncertainty rather than a low income level is associated with reduced cognitive test scores for farmers in Brazil, consistent with the interpretation that worry is an important driver of the effect.[31] Kaur, Mullainathan, et al. (2021) demonstrate that financial constraints can also impair worker productivity in a recent field experiment which varies whether workers are paid at the end of the working period or receive parts of their earnings upfront. The findings of Apenbrink (2021) on the adverse effect of Ebola concern on cognitive function in the US can be seen as an extension of this literature from worrying about income to worrying about health.

One potential explanation for the discrepancies between my findings here and those in Apenbrink (2021) is that exposure to different media reports about a given level of epidemic activity—i.e., a pure manipulation of perception—induces a weaker emotional response than the variation in actual epidemic activity exploited in the analysis of the cognitive cost of US Ebola cases. Specifically, the shift in worry caused by the current treatment manipulation might not be extreme enough to meaningfully affect subjects' cognitive function.[32] Presumably, the difference in the strength of emotional responses is magnified by (i) the subjectively more horrific nature and case fatality rate of Ebola compared to COVID-19, (ii) the difference in sample composition between young university students and old people from all social classes, and (iii) the later timing of the current experiment in the life cycle of the epidemic, when many people have probably adapted to the threat and the general level of uncertainty is low. In contrast, the analysis in Apenbrink (2021) explores a situation of high uncertainty just at the onset of a potential epidemic. From this perspective, the current study tests the limits of a meaningful effect of epidemic-induced worry on cognitive function.

An alternative interpretation is suggested by my findings regarding subjects' selective attention to pandemic-related news in Section 3.4.5 and their self-perceived productivity in Section 3.4.6. These pieces of evidence indicate that the emotional response evoked in the context

31. See de Bruijn and Antonides (2022) for a detailed review of the relevant studies and a discussion of the general state of the evidence.

32. Kaur, Mullainathan, et al. (2021) make a similar point about the difference between priming and alleviating financial strain in experimental tests of the causal effects of scarcity. This would also be consistent with the results of Bogliacino, Codagnone, Montealegre, Folkvord, Gómez, et al. (2021), who document negative associations between various self-reported adverse real-life shocks during the COVID-19 pandemic and cognitive function outcomes, but do not find an effect with an experimental priming intervention.

of the current experiment affects cognition, but the cognitive effects do not translate into meaningful differences in labor productivity. In contrast to pure tests of cognitive ability, solving mental arithmetic problems also requires a substantial degree of effort. Therefore, participants in the current study might be able to compensate for worry-induced reductions in task-available cognitive resources by putting in more mental effort, maybe with the intention of reaching an earnings goal that serves as their reference point.

Happiness and labor productivity

My study also connects to a large body of research on the relationship between employee well-being and employee productivity (reviewed, e.g., in Krekel, Ward, and De Neve, 2019). Much of this literature is correlational or unsuitable to pin down a direct effect of positive affect on productivity, but the findings from causal studies generally suggest a task-specific effect. Whereas Bellet, De Neve, and Ward (2020) find a negative effect of weather-induced reductions in happiness at work on the productivity of call center salespersons, who rely a lot on socioemotional skills, Borowiecki (2017) documents an increase in creative output for composers after unexpected family bereavements.

The one study that is similar enough to warrant closer comparison is Oswald, Proto, and Sgroi (2015), who investigate the causal effect of happiness on labor productivity in a series of laboratory experiments. Though the thematic focus is different, the design of my experiment shares many features with their setup. Both studies employ student samples from European universities, measure productivity on a mental arithmetic task under piece-rate incentives, and induce similarly strong exogenous changes in happiness.[33] In addition, two of the four experiments reported in Oswald, Proto, and Sgroi (2015) also use videos to manipulate emotions.[34] However, while they document an increase in productivity of about 12 percent, the coefficient estimates in my main specification are smaller in absolute value by a factor of ten. The lower bound of the 95 percent confidence interval of my estimate is at approximately −5.8 percent, implying that I can comfortably rule out a change in productivity that is half the size of that found by Oswald, Proto, and Sgroi (2015).

I see two starting points for further research to investigate the reasons for this discrepancy. First, a subtle difference in the design of my experiment is that I ask subjects for their goal number of correct answers before they start working on the mental arithmetic problems of the second session. This design feature is directly related to one potential explanation for my null result: if participants perceive falling short of their stated goal as a loss, this prospect could motivate them to exert the necessary extra effort to compensate for the cognitive effects of a change in happiness. Put differently, setting a goal might make exerted effort decline less in

33. Oswald, Proto, and Sgroi (2015) report a posttreatment level difference in happiness of 0.67 points on a 7-point scale in their second experiment, which also uses video clips as the source of exogenous variation. By comparison, my treatment manipulation induces a happiness difference of 1.45 points on an 11-point scale. Both estimates correspond to an effect size of about 0.7 SD.

34. In the other two experiments, they provide a random subset of participants with snacks and drinks at the beginning of the session or analyze the effect of natural variation in the experience of a major real-world shock like the recent death of a close family member.

response to increases in effort costs. The possibility that self-chosen goals can preserve productivity in the face of obstacles is also suggested by Kaur, Kremer, and Mullainathan (2015) and Clark, Gill, Prowse, and Rush (2020) in the context of self-control problems.[35]

Second and relatedly, the productivity consequences of changes in happiness might be asymmetric: happiness decreases might not have the same effect on productivity as comparable increases in happiness. In particular, it seems less likely that individuals will adjust their cognitive effort downwards if a boost to their well-being frees up additional cognitive resources by diverting their attention from pre-existing concerns. While Oswald, Proto, and Sgroi (2015)—in their fourth experiment—report a significant association between participants' productivity in the laboratory and natural variation in happiness after the experience of family tragedies that would cast doubt on this conjecture, they also acknowledge that this particular piece of evidence is not as convincing as their other experiments. For instance, it could also be explained by a correlation between productivity and latent health risks or life circumstances which increase the likelihood of experiencing illness or death.

3.6 Conclusion

By means of an online experiment, I show that negative emotional responses to the COVID-19 pandemic induced by media reports do not have a meaningful negative effect on labor productivity in a cognitively demanding task in a sample of university students eight months into the pandemic. This null result is robust across a variety of analyses and cannot be explained by unintended countervailing motivational effects of the treatment manipulation or cheating on the task. Yet, I provide suggestive evidence that worrying leads to a selective shift of attention towards pandemic-related news and an increase in the occurrence of distracting thoughts that make it harder for subjects to keep concentrated while working, in line with the intuition that negative emotions have adverse cognitive effects. One plausible interpretation of these findings is that subjects make up for worry-induced distractions by putting in more cognitive effort, consistent with the notion of income targeting.

My findings have implications for optimal public communication during epidemics: they indicate that exposure to information about the danger of the disease does not have a direct adverse effect on labor productivity. This is good news for policymakers who want to promote adherence to COVID-19 preventive measures by reiterating the threat of the virus in public communication. However, my findings also imply that this can come at the cost of short-term reductions in happiness above and beyond foregone utility from induced changes in behavior.

If one feels comfortable to extrapolate from my results to a comparison with the counterfactual situation without COVID-19, my findings also suggest that labor productivity declines of otherwise unaffected workers as a result of negative emotional responses are not a key channel by which epidemics disrupt the economy, at least not in later stages of the epidemic when people have had time to adapt to the new situation. However, in light of the context of my experiment and the pattern of findings in previous research on the effect of emotions on productivity,

35. Economic experiments also provide evidence that personal goals increase productivity more generally (e.g., Goerg and Kube, 2012).

it is unclear whether my conclusions readily extend to this comparison, which implies much stronger emotional reactions. Moreover, it is important to keep in mind that epidemic-induced worry can also harm the economy in other ways, e.g., by inducing deliberate behavioral changes of workers and consumers, and that it presumably comes with negative health consequences by itself (e.g., Blix, Birkeland, and Thoresen, 2021). Therefore, it still seems warranted to prevent the spread of excessive fears that may arise from an epidemic.

Future research should seek to clarify the conditions under which emotions affect economic outcomes. Of particular interest could be the role of salient goals, which might be an important factor for maintaining productivity in the face of worry and unhappiness.

Appendix 3.A Additional Figures and Tables

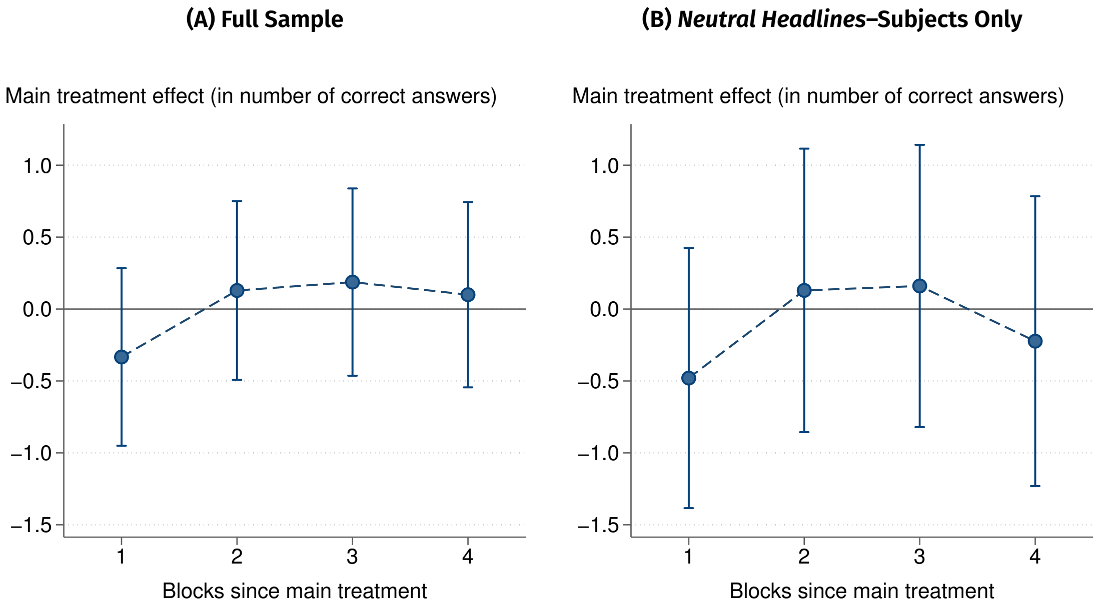

Figure 3.A.1. Evolution of the Main Treatment Effect over Time

Notes: Coefficient estimate plots of block-specific main treatment effects for each second-session block of mental arithmetic problems. Blocks 1 and 2 are non-news blocks, blocks 3 and 4 are news blocks. The estimates are from a regression of the subject-block-level number of correct answers on indicators for each second-session block and their interactions with an indicator for the *Worry-Amplifying* condition, controlling for the number of correct answers at baseline and the baseline level of worry. Error bars indicate 95 percent confidence intervals, constructed using standard errors that are robust to heteroscedasticity and arbitrary intra-cluster correlation within subjects. Panel (A) shows coefficient estimates from the full sample, Panel (B) uses only the subsample of subjects in the *Neutral Headlines* subcondition.

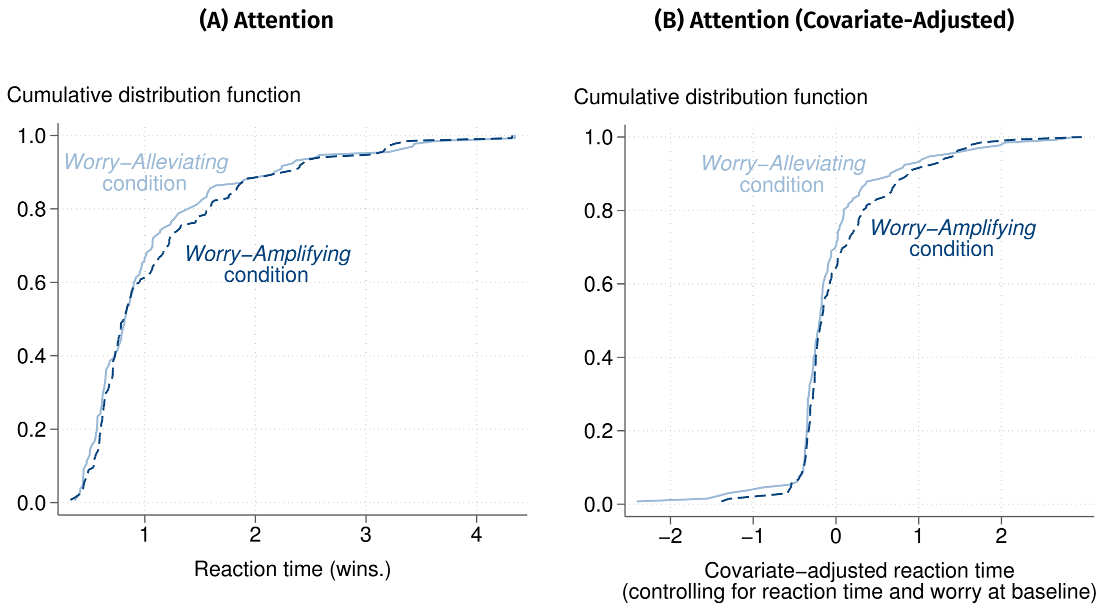

Figure 3.A.2. Attention in Non-news Blocks by Experimental Condition

Notes: The cumulative distribution of a measure of attention in non-news blocks, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of average reaction time in seconds across all mental arithmetic problems attempted in the two blocks. Panel (B) shows the distribution of the residuals from a regression of average reaction time on average reaction time at baseline and the baseline level of worry. Both reaction time variables used to calculate the cumulative distribution functions are winsorized by replacing the highest with the second-highest value in each condition.

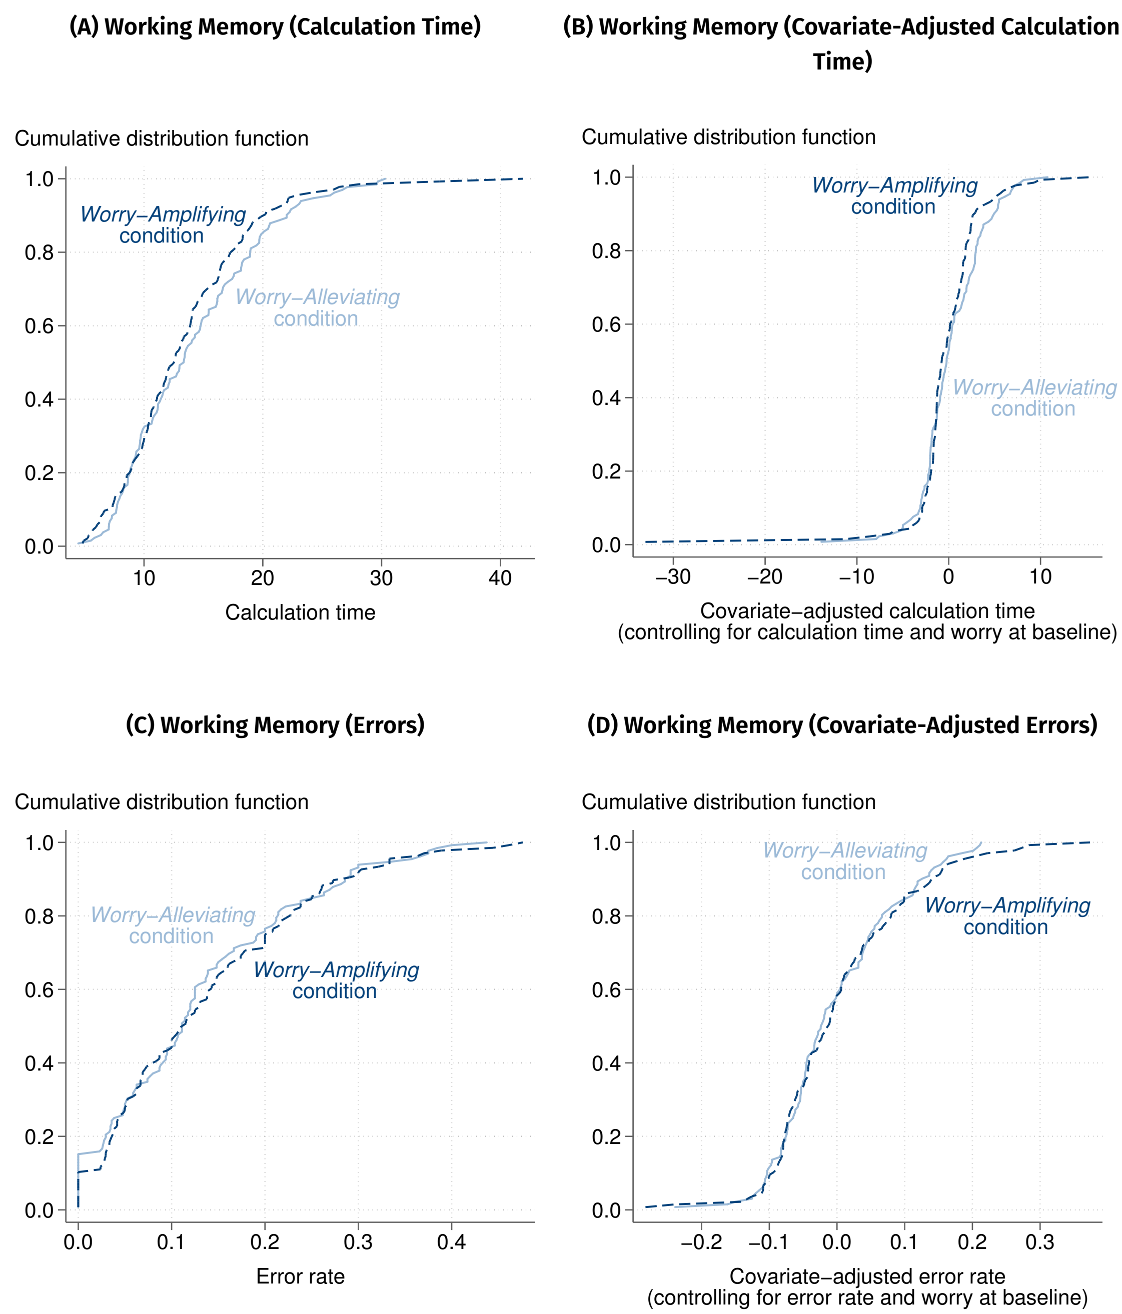

Figure 3.A.3. Working Memory in Non-news Blocks by Experimental Condition

Notes: The cumulative distribution of two measures of working memory in non-news blocks, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of average calculation time in seconds across all mental arithmetic problems attempted in the two blocks. Panel (B) shows the distribution of the residuals from a regression of average calculation time on average calculation time at baseline and the baseline level of worry. Panel (C) shows the distribution of the average error rate (i.e., the fraction of incorrect answers) across the two blocks. Panel (D) shows the distribution of the residuals from a regression of the average error rate on the baseline error rate and the baseline level of worry.

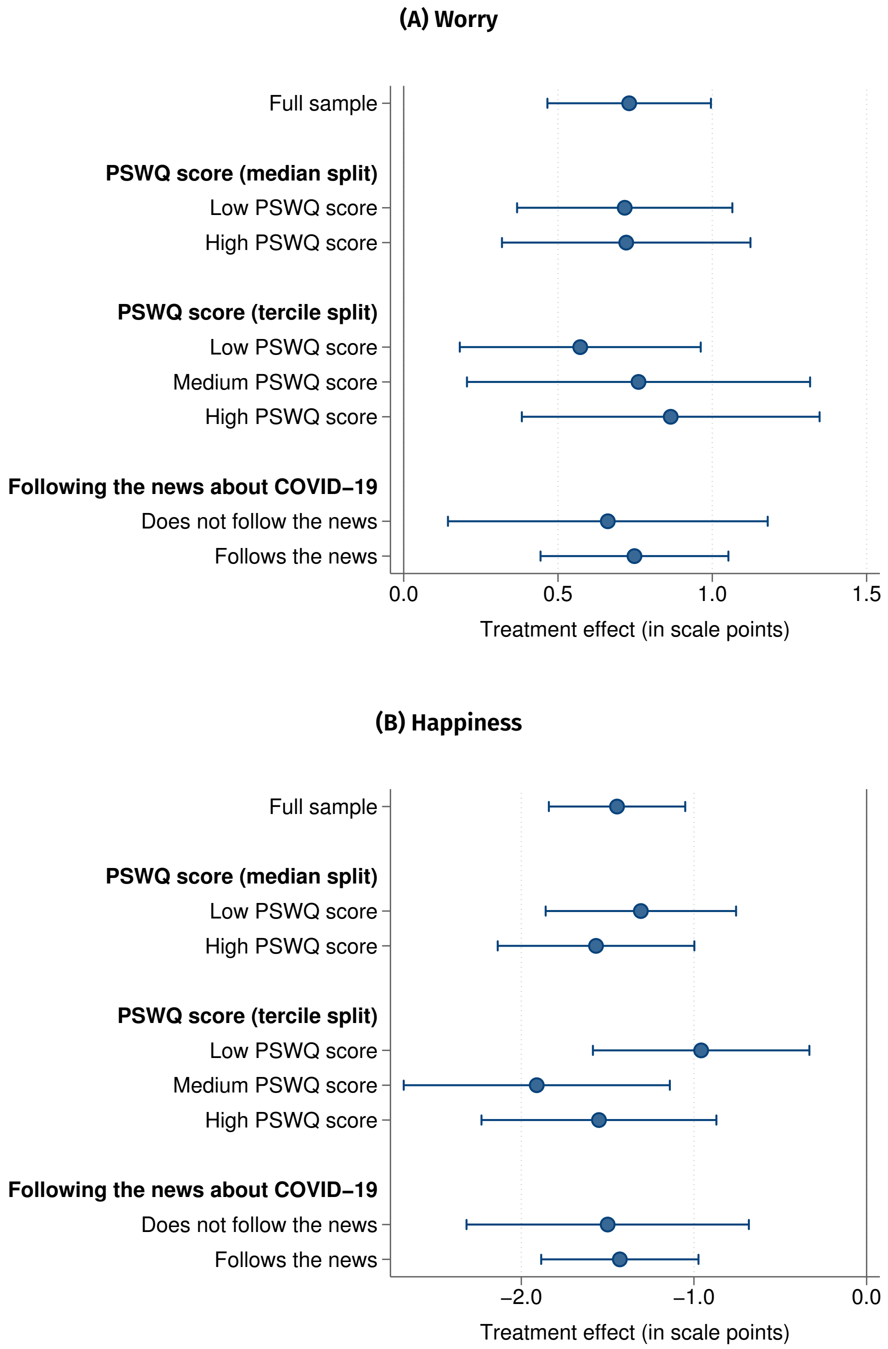

Figure 3.A.4. Heterogeneity in the Effect of the Main Treatment Manipulation on Worry and Happiness

Notes: Plots of coefficient estimates of the effect of the main treatment manipulation on worry and happiness for the full sample and three different partitions of the full sample into subsamples: a median split by PSWQ score, a tercile split by PSWQ score, and a split by whether or not subjects report following the news about COVID-19. The estimates are from regressions of subjects' level of worry or happiness after the main treatment manipulation on indicators for each of that partition's subsamples and their interactions with an indicator for the *Worry-Amplifying* condition, controlling for the level of worry or happiness at baseline. Error bars indicate 95 percent confidence intervals, constructed using heteroscedasticity-robust standard errors. Panel (A) shows coefficient estimates for the treatment effect on worry, Panel (B) for the effect on happiness.

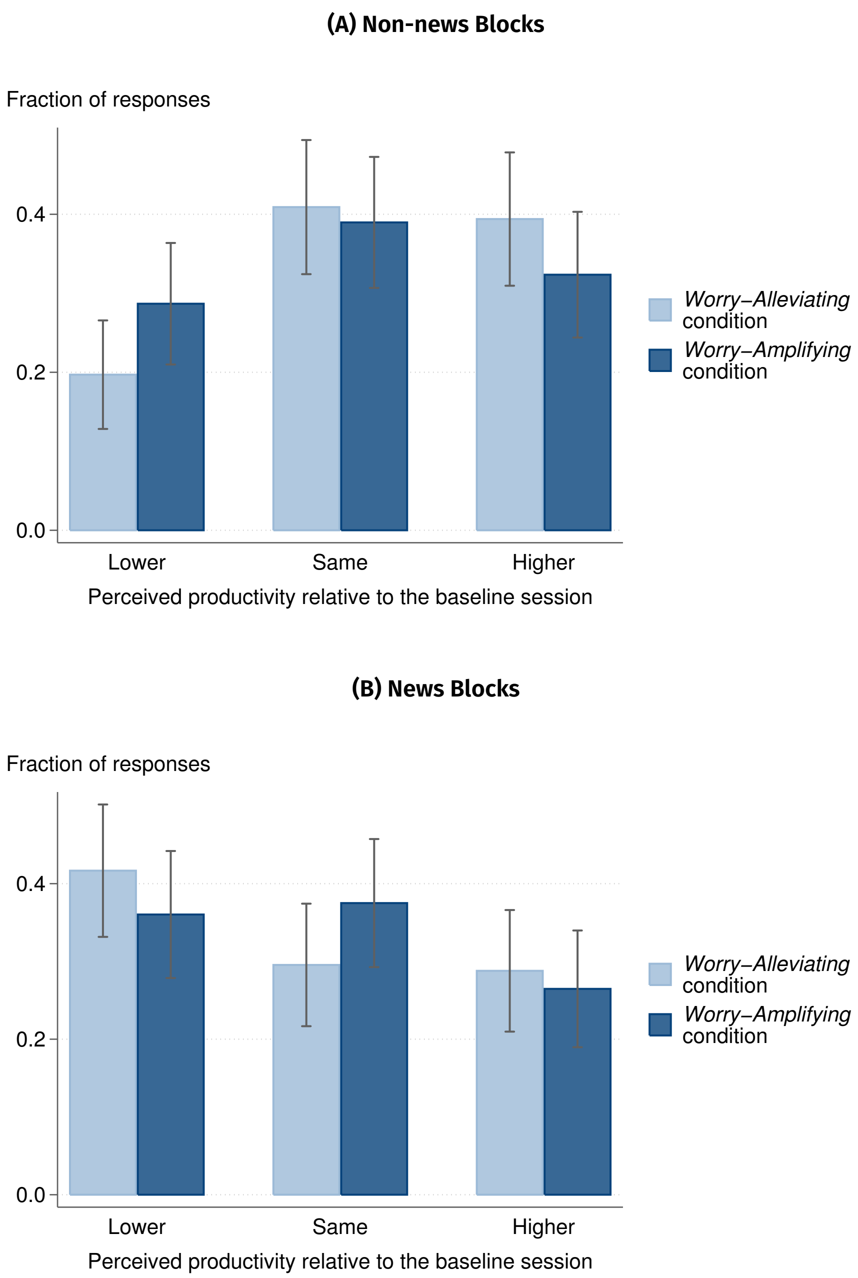

Figure 3.A.5. Perceived Productivity in Non-news Blocks and News Blocks by Experimental Condition

Notes: Bar charts of productivity perceptions for the two types of second-session blocks relative to the baseline session, plotted separately for subjects in the *Worry-Alleviating* condition (in light blue) and the *Worry-Amplifying* condition (in dark blue). Panel (A) shows the distribution of responses for relative productivity in non-news blocks, Panel (B) for relative productivity in news blocks.

Table 3.A.1. Descriptive Statistics for the Analysis Sample

Variable	Mean	SD	Min.	Median	Max.	Obs.
Worry	6.425	2.327	0	7	10	268
Worry at baseline	6.093	2.369	0	7	10	268
Change in worry	0.332	1.235	−3	0	7	268
Happiness	5.817	2.134	0	6	10	268
Happiness at baseline	6.224	1.972	0	7	10	268
Change in happiness	−0.407	1.972	−7	0	8	268
Female	0.604	0.490	0	1	1	268
Age	25.922	7.191	16	24	72	268
Bachelor's degree	0.321	0.468	0	0	1	268
Master's degree	0.190	0.393	0	0	1	268
High school GPA	2.168	0.751	1.150	2.150	5.000	268
No high school GPA	0.022	0.148	0	0	1	268
High school math grade	2.298	0.902	1.150	2.150	5.000	268
No high school math grade	0.007	0.086	0	0	1	268
Correct answers at baseline	10.959	3.952	1	10	21	268
Correct answers per non-news block	12.058	3.889	3.500	11.500	21.500	268
Change in productivity	1.099	2.404	−8.000	1.000	12.000	268
Correct answers per news block	12.326	4.024	3.000	12.250	24.000	268
Reaction time at baseline	1.864	9.639	0.307	0.916	157.706	267
Reaction time at baseline (wins.)	1.293	1.141	0.307	0.916	7.550	267
Reaction time in non-news blocks	1.191	1.629	0.328	0.818	24.357	267
Reaction time in non-news blocks (wins.)	1.110	0.789	0.328	0.818	4.345	267
Calculation time at baseline	15.029	6.792	3.673	13.961	68.660	267
Calculation time in non-news blocks	13.749	5.750	4.464	12.936	41.889	267
Error rate at baseline	0.167	0.159	0.000	0.133	0.800	268
Error rate in non-news blocks	0.132	0.107	0.000	0.111	0.476	268
Attempted answers at baseline	12.959	3.427	1	13	23	268
Goal for correct answers per block	11.388	4.693	2	10	50	268
PSWQ score	48.172	13.090	16	46	80	268
Follows the news about COVID-19	0.761	0.427	0	1	1	268
Video contained new information	0.500	0.501	0	1	1	268
Belief about infection risk	27.560	23.753	0	20	100	268
Belief about risk of long-term effects	11.060	13.185	0	9	85	268
Belief about risk of death	3.575	6.689	0	1	60	268
Belief about months till vaccination	8.220	7.488	1	6	60	268
Does not believe in vaccination	0.004	0.061	0	0	1	268
Worry about financial situation	3.534	2.682	0	3	10	268
Worry about job prospects	3.869	2.902	0	4	10	268
Number of saved articles	1.866	2.828	0	0	15	268
At least one saved article	0.481	0.501	0	0	1	268
Satisfaction with own performance	6.698	2.115	0	7	10	268
Concentration on the task impaired by video	0.123	0.329	0	0	1	268

Notes: Mean, standard deviation, minimum, median, maximum, and number of observations of important variables for the analysis sample.

Table 3.A.2. Descriptive Statistics by Experimental Condition and Test of Balance

Variable	*Worry-Alleviating* condition	*Worry-Amplifying* condition	Difference of means	Standardized difference of means	Test of balance *p*-value
Worry at baseline	6.432	5.765	0.667	0.282	0.023
	(2.245)	(2.447)			
Happiness at baseline	6.189	6.257	−0.068	−0.034	0.967
	(2.112)	(1.834)			
Female	0.629	0.581	0.048	0.098	0.423
	(0.485)	(0.495)			
Age	25.879	25.963	−0.084	−0.012	0.911
	(7.349)	(7.062)			
Bachelor’s degree	0.318	0.324	−0.005	−0.011	0.925
	(0.468)	(0.470)			
Master’s degree	0.174	0.206	−0.032	−0.080	0.509
	(0.381)	(0.406)			
High school GPA	2.155	2.181	−0.026	−0.035	0.623
	(0.770)	(0.735)			
High school math grade	2.247	2.347	−0.100	−0.111	0.313
	(0.904)	(0.901)			
Correct answers at baseline	10.871	11.044	−0.173	−0.044	0.909
	(3.916)	(4.000)			
Reaction time at baseline	1.353	2.364	−1.011	−0.105	0.796
	(1.353)	(13.495)			
Reaction time at baseline (wins.)	1.349	1.238	0.110	0.097	0.801
	(1.329)	(0.923)			
Calculation time at baseline	15.098	14.962	0.136	0.020	0.460
	(5.703)	(7.731)			
Error rate at baseline	0.165	0.169	−0.004	−0.027	0.520
	(0.165)	(0.152)			
Attempted answers at baseline	12.788	13.125	−0.337	−0.098	0.346
	(3.288)	(3.562)			
PSWQ score	48.455	47.897	0.557	0.043	0.802
	(12.759)	(13.445)			
Follows the news about COVID-19	0.735	0.787	−0.052	−0.122	0.319
	(0.443)	(0.411)			
Observations	132	136			

Notes: Columns (1) and (2) show the mean values for selected (baseline) subject characteristics by experimental condition, with standard deviations in parentheses. Columns (3) and (4) show the simple and standardized difference of means between the two conditions, where the denominator of the standardized difference is the pooled-sample standard deviation. Column (5) shows the p-value of a test of balance across conditions, which is a Mann–Whitney U test for continuous and non-binary discrete variables and a Pearson’s χ^2 test for binary variables.

Table 3.A.3. Effect of the Main Treatment Manipulation on Measures of Attention

Dependent variable:	Reaction time in non-news blocks			Reaction time in non-news blocks (wins.)		
	(1)	(2)	(3)	(4)	(5)	(6)
Worry-Amplifying condition	0.191 (0.199)	0.166 (0.207)	0.200 (0.214)	0.053 (0.097)	0.097 (0.079)	0.102 (0.081)
Reaction time at baseline (cent.)		0.025 (0.389)	0.024 (0.383)			
Reaction time at baseline (wins., cent.)					0.396*** (0.070)	0.393*** (0.071)
Worry at baseline (cent.)			0.050 (0.038)			0.008 (0.014)
Constant	1.094*** (0.074)	1.107*** (0.212)	1.135*** (0.195)	1.083*** (0.069)	1.061*** (0.057)	1.066*** (0.056)
Observations	267	267	267	267	267	267
R^2 (adjusted)	0.000	0.018	0.019	−0.003	0.323	0.321
Mean (dependent variable)	1.191	1.191	1.191	1.110	1.110	1.110

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. The dependent variable is the average reaction time in seconds across all mental arithmetic problems attempted in non-news blocks. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. The abbreviation "wins." indicates replacement of the highest by the second-highest value of the variable in each condition. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.A.4. Effect of the Main Treatment Manipulation on Measures of Working Memory

Dependent variable:	Calculation time in non-news blocks			Error rate in non-news blocks		
	(1)	(2)	(3)	(4)	(5)	(6)
Worry-Amplifying condition	−0.691	−0.606	−0.502	0.006	0.004	0.006
	(0.706)	(0.542)	(0.535)	(0.013)	(0.012)	(0.012)
Calculation time at baseline (cent.)		0.619***	0.611***			
		(0.195)	(0.195)			
Error rate at baseline (cent.)					0.323***	0.322***
					(0.045)	(0.045)
Worry at baseline (cent.)			0.156*			0.002
			(0.083)			(0.002)
Constant	14.098***	14.055***	14.144***	0.129***	0.129***	0.130***
	(0.499)	(0.307)	(0.317)	(0.009)	(0.008)	(0.008)
Observations	267	267	267	268	268	268
R^2 (adjusted)	0.000	0.535	0.537	−0.003	0.222	0.220
Mean (dependent variable)	13.749	13.749	13.749	0.132	0.132	0.132

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. The dependent variable is the average calculation time in seconds or the average error rate (i.e., the fraction of incorrect answers) across all mental arithmetic problems attempted in non-news blocks. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.A.5. Matrix of Correlation Coefficients between Changes in Labor Productivity, Worry, and Happiness

Variable:	Change in worry	Change in happiness	Change in productivity
Change in worry	1.000		
Change in happiness	−0.165***	1.000	
	(0.007)		
Change in productivity	0.033	0.009	1.000
	(0.595)	(0.880)	

Notes: Matrix of Spearman's rank correlation coefficients, with two-sided p-values in parentheses. *Change in productivity* is the difference in the average number of correct answers per block between non-news blocks and baseline. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table 3.A.6. News Salience and Labor Productivity

	Dependent variable: Correct answers per news block		
	Worry-Alleviating subsample	*Worry-Amplifying* subsample	Full sample
	(1)	(2)	(3)
COVID-19 Headlines subcondition	−0.052	0.204	−0.055
	(0.399)	(0.462)	(0.401)
Worry-Amplifying condition			0.050
			(0.468)
Worry-Amplifying condition × *COVID-19 Headlines* subcondition			0.268
			(0.606)
Correct answers at baseline (cent.)	0.821***	0.800***	0.810***
	(0.050)	(0.055)	(0.038)
Worry at baseline (cent.)	0.113	−0.083	0.005
	(0.086)	(0.110)	(0.072)
Constant	11.541***	11.446***	11.491***
	(0.334)	(0.403)	(0.325)
Observations	132	136	268
R^2 (adjusted)	0.655	0.597	0.623
Mean (dependent variable)	12.167	12.482	12.326

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* and *COVID-19 Headlines subcondition* are indicators equal to one for subjects in the experimental conditions of the respective name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.A.7. News Salience and Labor Productivity (without Controlling for Baseline Worry)

	Dependent variable: Correct answers per news block		
	Worry-Alleviating subsample	*Worry-Amplifying* subsample	Full sample
	(1)	(2)	(3)
COVID-19 Headlines subcondition	−0.057	0.209	−0.055
	(0.399)	(0.460)	(0.400)
Worry-Amplifying condition			0.046
			(0.466)
Worry-Amplifying condition × *COVID-19 Headlines* subcondition			0.268
			(0.603)
Correct answers at baseline (cent.)	0.812***	0.807***	0.809***
	(0.051)	(0.054)	(0.037)
Constant	11.487***	11.539***	11.489***
	(0.322)	(0.368)	(0.318)
Observations	132	136	268
R^2 (adjusted)	0.653	0.598	0.624
Mean (dependent variable)	12.167	12.482	12.326

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* and *COVID-19 Headlines subcondition* are indicators equal to one for subjects in the experimental conditions of the respective name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Appendix 3.B Are Emotional Responses to the Videos Driven by Changes in Beliefs?

In section Section 3.4.2, I document that the news video clips displayed to subjects as part of the main treatment manipulation provoke a multi-faceted negative emotional response. Especially the worry-inducing video in the *Worry-Amplifying* condition significantly increases subjects' worry about the health consequences of COVID-19 and concurrently reduces their happiness. This could either be driven by new factual information in the videos or by the conveyed feelings and emotions. I conduct two analyses with the aim to distinguish these two forces.

First, I make use of participants' self-reports about the information content of their condition's video clip. Descriptively, exactly 50 percent of subjects report in the postexperimental questionnaire that the video they have seen provided them with new information.[36] The proportion of subjects who report seeing new information is nine percentage points lower in the *Worry-Amplifying* relative to the *Worry-Alleviating* condition, but this difference is not statistically significant ($p = 0.143$ in a two-sided z-test of proportions). Thus, the information channel is only plausible for about half the sample, and even slightly less in the condition that drives the emotional response.

Second, I test whether the main treatment manipulation has a lasting effect on stated beliefs about COVID-19. In Table 3.B.1 and Table 3.B.2, I report estimates of the effect of watching the fear-evoking video on beliefs about the risk of a COVID-19 infection in Germany within six months of the experiment, the risk of long-term effects after an infection, the risk of death due to an infection, and the number of months until a vaccine will be available to a majority of the German population.[37]

In line with the provision of new information in the videos, the coefficient estimates indicate sizable belief effects of up to about 20 to 25 percent relative to the mean of the *Worry-Alleviating* condition for the risk of long-term effects and the number of months till vaccination. However, due to the large amount of heterogeneity in beliefs, only the effect on beliefs about the wait time until a vaccine becomes available is consistently significant at the ten percent level.[38] Moreover, since subjects in the *Worry-Alleviating* condition overestimate the risk of long-term health effects after a COVID-19 infection on average, the information provided in the fear-inducing video reduced beliefs about this event.

Taken together, the observed pattern of belief changes does not fit the documented emotional responses to the news video clips, suggesting that their effects are not primarily driven by information provision.

36. All participants but one report remembering what the video is about.

37. Corresponding empirical distribution functions of the outcome variables by main experimental condition are presented in Figure 3.B.1 to Figure 3.B.4 below.

38. One subject reported the belief that no vaccine against COVID-19 would ever be available. That subject's belief was coded as the highest number of months observed in the data, i.e., 60 months. The effect continues to be significant at the ten percent level if that subject is dropped from the sample instead.

Table 3.B.1. Effect of the Main Treatment Manipulation on Beliefs about the Pandemic

Dependent variable:	Belief about infection risk		Belief about risk of long-term effects	
	(1)	(2)	(3)	(4)
Worry-Amplifying condition	3.059	3.970	−3.226**	−2.616
	(2.910)	(2.971)	(1.614)	(1.689)
Worry at baseline (cent.)		1.366**		0.915***
		(0.606)		(0.334)
Constant	26.008***	26.784***	12.697***	13.217***
	(2.009)	(2.012)	(1.275)	(1.282)
Observations	268	268	268	268
R^2 (adjusted)	0.000	0.015	0.011	0.034
Mean (dependent variable)	27.560	27.560	11.060	11.060

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. The dependent variable is measured in the number of individuals out of 100 who will suffer the respective outcome in Germany within six months of the experiment, unconditional in column (1) and conditional on contracting COVID-19 in column (2). *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.B.2. Effect of the Main Treatment Manipulation on Beliefs about the Pandemic (II)

Dependent variable:	Belief about risk of death		Belief about months till vaccination	
	(1)	(2)	(3)	(4)
Worry-Amplifying condition	−0.032	0.271	1.688*	1.781*
	(0.822)	(0.835)	(0.912)	(0.914)
Worry at baseline (cent.)		0.455***		0.140
		(0.157)		(0.208)
Constant	3.591***	3.849***	7.364***	7.443***
	(0.599)	(0.632)	(0.604)	(0.612)
Observations	268	268	268	268
R^2 (adjusted)	−0.004	0.018	0.009	0.007
Mean (dependent variable)	3.575	3.575	8.220	8.220

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. The dependent variable is measured in the number of infected individuals out of 100 who will die as a consequence of the disease in Germany for columns (1) and (2), and in the number of months until a vaccine will become available for the majority of the German population in columns (3) and (4). *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. In columns (3) and (4), the dependent variable is coded as the maximum observed belief in the data (60 months) for one subject who does not believe that a vaccine will ever become available. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

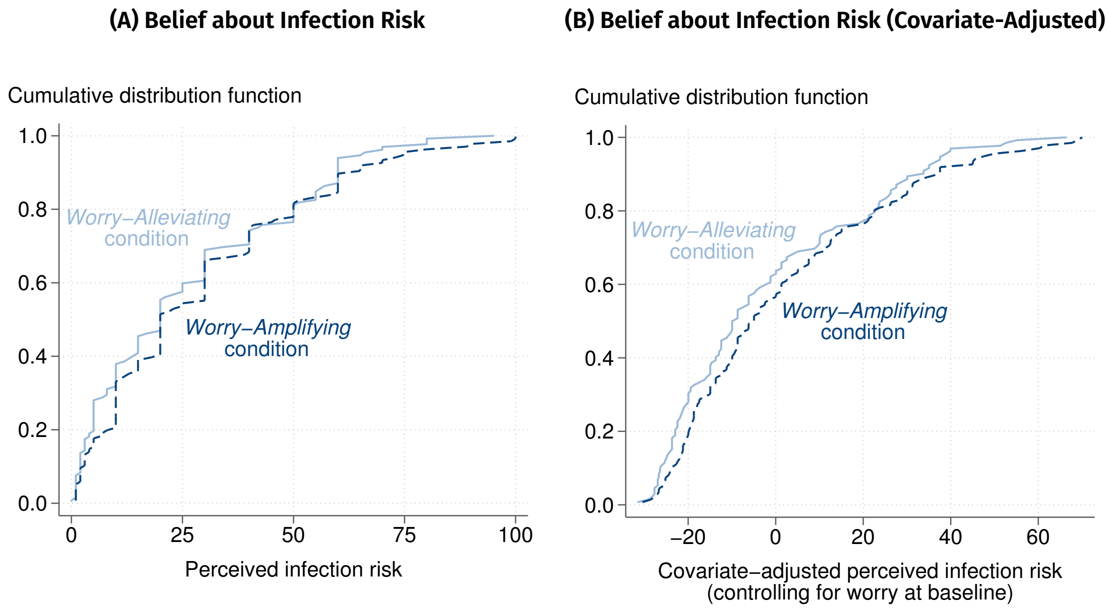

Figure 3.B.1. Beliefs about COVID-19 Infection Risk by Experimental Condition

Notes: The cumulative distribution of beliefs about the risk of a COVID-19 infection in Germany within six months of the experiment, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of perceived infection risk in percent. Panel (B) shows the distribution of the residuals from a regression of perceived infection risk on the level of worry at baseline.

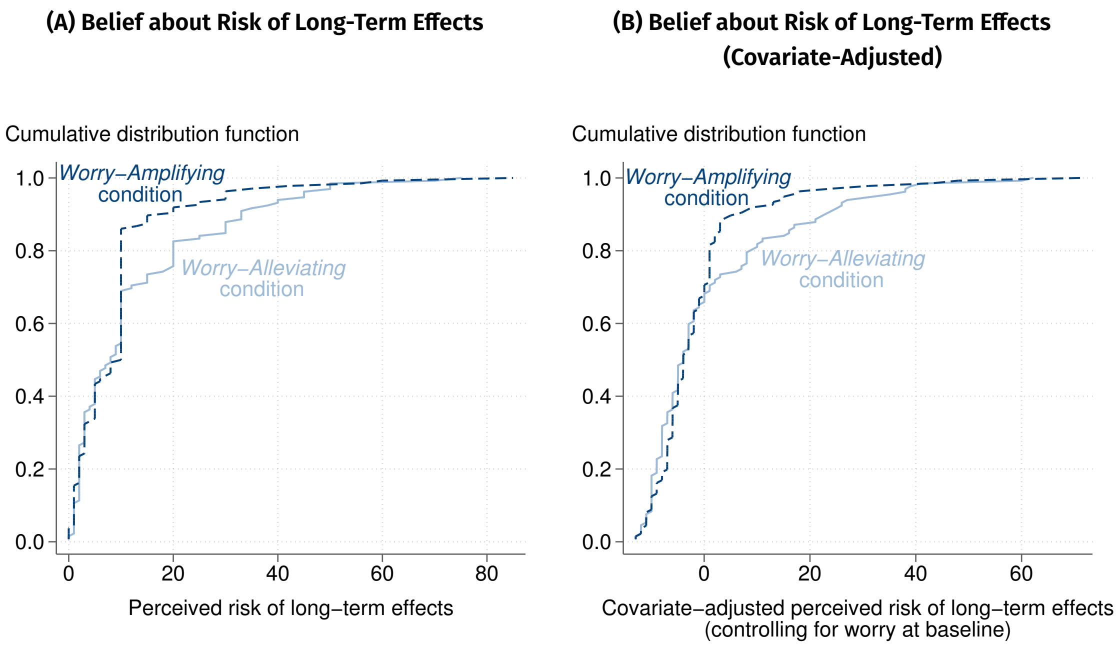

Figure 3.B.2. Beliefs about the Risk of Long-Term Effects after a COVID-19 Infection by Experimental Condition

Notes: The cumulative distribution of beliefs about the risk that a COVID-19 infection causes long-term effects in Germany, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of perceived risk of long-term effects conditional on an infection in percent. Panel (B) shows the distribution of the residuals from a regression of the perceived risk of long-term effects on the level of worry at baseline.

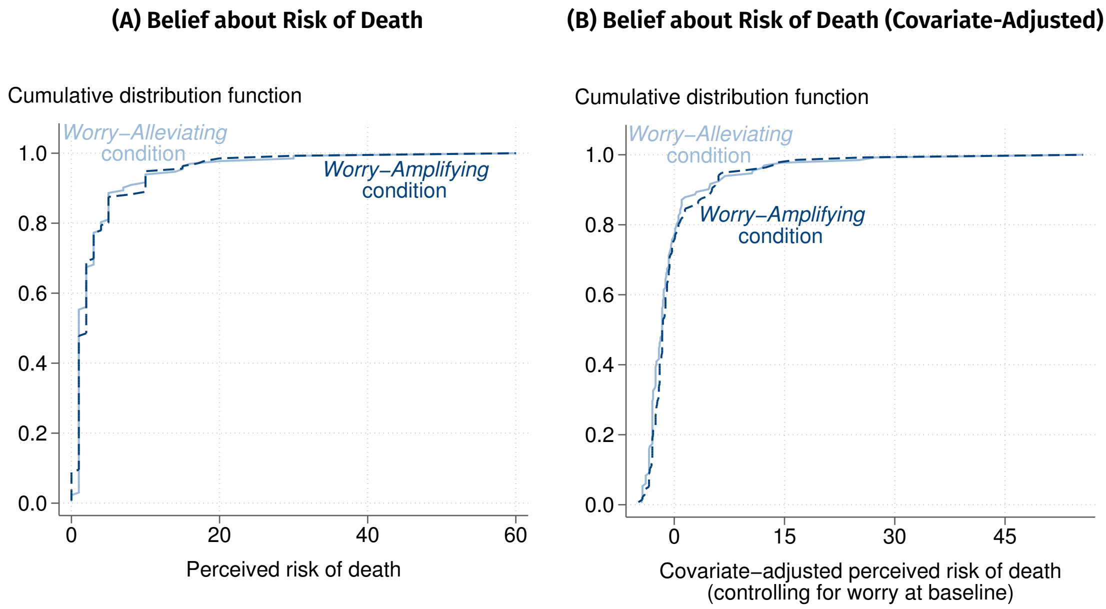

Figure 3.B.3. Beliefs about the Risk of Death Due to a COVID-19 Infection by Experimental Condition

Notes: The cumulative distribution of beliefs about the risk of death due to a COVID-19 infection in Germany, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). Panel (A) shows the distribution of perceived risk of death conditional on an infection in percent. Panel (B) shows the distribution of the residuals from a regression of perceived risk of death on the level of worry at baseline.

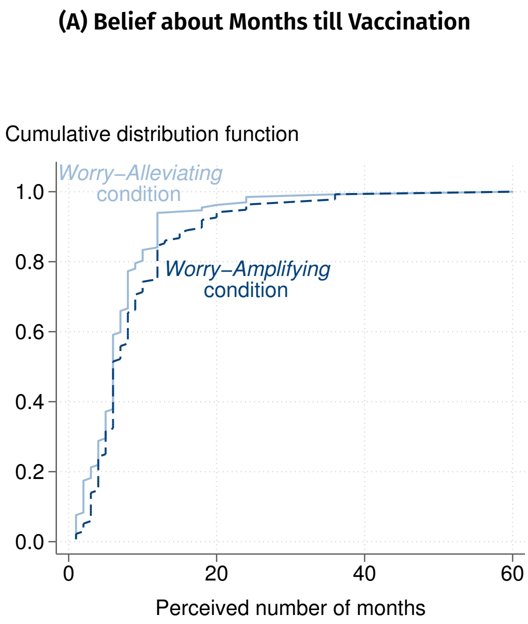

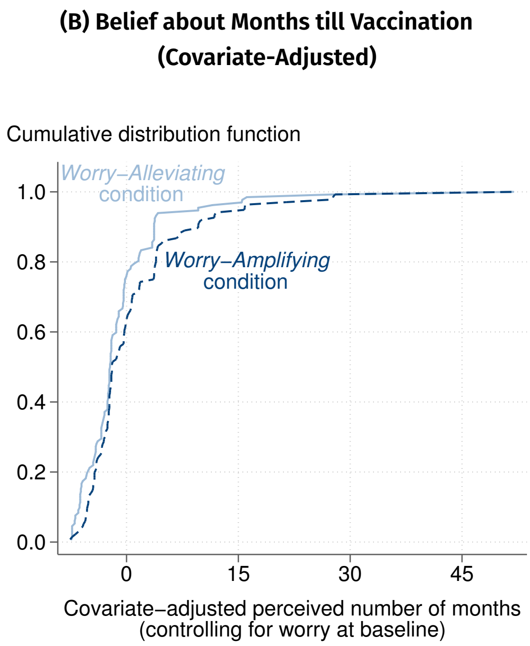

Figure 3.B.4. Beliefs about the Number of Months until COVID-19 Vaccination by Experimental Condition

Notes: The cumulative distribution of beliefs about the number of months until a COVID-19 vaccine will be available for the majority of the German population, plotted separately for subjects in the *Worry-Amplifying* condition (dashed dark blue line) and the *Worry-Alleviating* condition (solid light blue line). The variable is coded as the maximum observed belief in the data (60 months) for one subject who does not believe that a vaccine will ever become available. Panel (A) shows the distribution of the perceived number of months till vaccination. Panel (B) shows the distribution of the residuals from a regression of the perceived number of months till vaccination on the level of worry at baseline.

Appendix 3.C Robustness

Table 3.C.1. Effect of the Main Treatment Manipulation on Labor Productivity in Non-news Blocks (Alternative Control Variables)

	Dependent variable: Correct answers per non-news block				
	(1)	(2)	(3)	(4)	(5)
Worry-Amplifying condition	−0.146	−0.185	−0.138	−0.267	−0.268
	(0.288)	(0.285)	(0.287)	(0.302)	(0.312)
Correct answers at baseline (cent.)	0.787***	0.780***	0.790***	0.783***	0.772***
	(0.036)	(0.035)	(0.033)	(0.034)	(0.038)
Worry at baseline (cent.)	−0.112		−0.104		
	(0.070)		(0.064)		
Happiness at baseline (cent.)			0.018		
			(0.071)		
Constant	11.291***	10.945***	11.290***	10.612***	10.594***
	(0.275)	(0.364)	(0.206)	(0.446)	(0.515)
Demographic controls	Yes	No	No	No	Yes
Baseline worry dummies	No	Yes	No	Yes	Yes
Baseline happiness dummies	No	No	No	Yes	Yes
Observations	268	268	268	268	268
R^2 (adjusted)	0.653	0.669	0.658	0.666	0.660
Mean (dependent variable)	12.058	12.058	12.058	12.058	12.058

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. *Worry-Amplifying condition* is an indicator equal to one for subjects in the experimental condition of the same name. Demographic controls include high school GPA, high school math grade, and dummies for gender, 10-year age group, highest obtained university degree (Bachelor's or Master's degree), and not reporting a GPA or math grade. Subjects who don't report a high school grade are assigned a value of 5.0, i.e., the worst possible grade. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Table 3.C.2. Main Results (excluding Suspected Cheaters)

Dependent variable:	Worry	Happiness	Correct answers per non-news block	Correct answers per news block
	(1)	(2)	(3)	(4)
Worry-Amplifying condition	0.711***	−1.356***	−0.186	−0.038
	(0.133)	(0.213)	(0.271)	(0.421)
COVID-19 Headlines subcondition				0.174
				(0.385)
Worry-Amplifying condition × *COVID-19 Headlines* subcondition				0.193
				(0.581)
Correct answers at baseline (cent.)			0.818***	0.856***
			(0.035)	(0.037)
Worry at baseline (cent.)	0.874***		−0.152***	−0.004
	(0.029)		(0.057)	(0.063)
Happiness at baseline (cent.)		0.600***		
		(0.067)		
Constant	6.800***	6.984***	11.040***	11.148***
	(0.081)	(0.122)	(0.188)	(0.282)
Observations	247	247	247	247
R^2 (adjusted)	0.783	0.401	0.697	0.676
Mean (dependent variable)	6.393	5.830	11.905	12.128

Notes: OLS estimates, with heteroscedasticity-robust standard errors in parentheses. 21 subjects are excluded for suspected cheating on the task according to the criteria described in Section 3.5.1. *Worry-Amplifying condition* and *COVID-19 Headlines subcondition* are indicators equal to one for subjects in the experimental conditions of the respective name. * denotes $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$, all for two-sided hypothesis tests.

Appendix 3.D Experimental Instructions

This appendix reproduces the instructions of the experiment, translated from the German original. Differences between experimental conditions and explanatory remarks about specific features of the instructions or the experimental interface are indicated in brackets. For legal reasons, the elicitation of subjects' bank details was conducted at the end of the second session rather than at the beginning of the first session for subjects who participated in the second and third iteration of the experiment during the first week of December.

3.D.1 Instructions for Session 1

Screen 1—Welcome to the online study of the BonnEconLab.

[introductory screen with informed consent form and general information about the dates and duration of the two sessions and the payment procedure]

Screen 2—Your bank details.

[form to elicit subjects' bank details]

Screen 3—Questionnaire.

Before the start of the study, please answer the following questions:

What is your gender?

How old are you?

What is your highest level of education?

What is your current or last major?

What was your GPA in high school?

What was your last grade in math class in high school?

How would you rate your happiness at the moment? Please use a scale from 0 to 10, where 0 means "not happy at all" and 10 means "very happy".

How worried are you that COVID-19 could cause serious damage to your health or the health of your loved ones? Please use a scale from 0 to 10, where 0 means "not worried at all" and 10 means "very worried".

Screen 4—Instructions for session 1 of the study.

In this session, you can earn a payoff by solving mental arithmetic problems. You will receive €0.10 for every correctly solved problem.
There will be two blocks of 5 minutes each, during which you can solve as many problems as you like. Technically, there is a maximal number of problems, but you will not reach this limit when making a serious attempt at solving the problems.
Note: Please don't refresh the page or switch to other tabs during the mental arithmetic problems.

The sequence of events for a problem:

The sequence of events for each problem is as follows:

- First, you will see an empty white screen.

- After a few seconds, a blurry text will appear. From that point on, you can display the mental arithmetic problem by simultaneously holding down the two keys Q and ENTER (and no other key) on your keyboard. You will then see an arithmetic problem consisting of four double-digit numbers, some of which have to be added up and others subtracted.
- You are supposed to solve the displayed problem by means of mental arithmetic. The problem will be displayed as long as you keep holding down the two keys Q and ENTER.
- As soon as you release the keys Q and ENTER or press any other additional key, the problem will disappear and be replaced by an input field. You will then have **exactly 5 seconds** to type in your answer. Please only enter numbers into the input field.
- After the 5 seconds have elapsed, the input field will disappear **automatically** and the sequence of events for the next problem begins.

Note: Make sure to not accidentally activate the caps lock key ⇩ during the mental arithmetic problems.

Solution of an example problem:

Imagine that one of the mental arithmetic problems is: **65 + 11 – 37 + 29**. The correct solution to this problem is **68**. If this answer was in the input field after the 5 seconds have elapsed, you would receive a payoff of €0.10 for this problem.

You would *not* receive a payoff for this problem in the following two cases:

- if some other number was in the input field after the 5 seconds have elapsed;
- if the input field disappeared after the end of the 5 seconds before you had typed in the correct solution.

Screen 5—Comprehension questions for session 1 of the study.

Question 1:

The purpose of this question is to test whether your device meets the technical requirements for the mental arithmetic problems. There is a blurry text below this question. A number should appear in place of the blurry text when you hold down the two keys Q and ENTER (and no other key) on your keyboard. What number is displayed in place of the blurry text when you hold down the keys Q and ENTER?

Question 2:

Imagine that one of the mental arithmetic problems in session 1 of the study is: **31 – 12 + 19 + 53**. What is the correct solution to this problem, for which you would receive a payoff of €0.10?

Question 3:

What would be your total payoff from the first session of the study if you answered 18 mental arithmetic problems correctly and 6 problems incorrectly during the two blocks?

◯ €0.00
◯ €1.80
◯ €2.40
◯ impossible to tell

Screen 6.

Block 1 will start in a few seconds.

Screen 7.
[first 5-minute block of mental arithmetic problems]

Screen 8—Break.
[a relaxing image is displayed for 30 seconds]

Screen 9.
Block 2 will start in a few seconds.

Screen 10.
[second 5-minute block of mental arithmetic problems]

Screen 11—End of the first session of the study.
You have successfully completed the first session of the study.
The link to start the second session of the study will be sent to you via e-mail **tomorrow, [date],** at about 6:15. The second session will have to be completed until 23:59.

3.D.2 Instructions for Session 2

Screen 1—Welcome to the second session of the study.
[introductory screen]

Screen 2—Instructions for session 2 of the study.
This session of the study consists of two parts.
Information about part 1:
In part 1, you will watch a short media report about a current issue. Afterwards, there will be a question about the content of the media report. You will have **60 seconds** to answer the question.
For this part, you will receive a payoff of €4.00.
Information about part 2:
In part 2, you can once again earn a payoff by solving mental arithmetic problems. The rules will be the same as in session 1 of the study. You will receive €0.10 for every correctly solved problem.
There will be four two blocks of 5 minutes each, during which you can solve as many problems as you like.
Note: Please don't refresh the page or switch to other tabs during the mental arithmetic problems.
The sequence of events for a problem in part 2:
The sequence of events for each problem is as follows:
- First, you will see an empty white screen.
- After a few seconds, a blurry text will appear. From that point on, you can display the mental arithmetic problem by simultaneously holding down the two keys Q and ENTER (and no other key) on your keyboard. You will then see an arithmetic problem consisting of four double-digit numbers, some of which have to be added up and others subtracted.

- You are supposed to solve the displayed problem by means of mental arithmetic. The problem will be displayed as long as you keep holding down the two keys Q and ENTER.
- As soon as you release the keys Q and ENTER or press any other additional key, the problem will disappear and be replaced by an input field. You will then have **exactly 5 seconds** to type in your answer. Please only enter numbers into the input field.
- After the 5 seconds have elapsed, the input field will disappear **automatically** and the sequence of events for the next problem begins.

Note: Make sure to not accidentally activate the caps lock key ⇩ during the mental arithmetic problems.

Special feature in part 2:

Sometimes, headlines of current news articles will be displayed to you for about 10 seconds in between two mental arithmetic problems. Meanwhile, the timer for solving problems will be paused. Whether and when you will see headlines doesn't depend on your decisions.
If you want to read one of the displayed news articles, you can bookmark it for later by clicking the button "Save this news article". Bookmarked news articles will be redisplayed to you **at the end of the study**.

Screen 3—Comprehension questions for session 2 of the study.

Question 1:

This question tests whether you can hear played sounds with your device. Please play back the following audio file. What number is stated in the recording?

Question 2:

The purpose of this question is to test whether your device meets the technical requirements for the mental arithmetic problems. There is a blurry text below this question. A number should appear in place of the blurry text when you hold down the two keys Q and ENTER (and no other key) on your keyboard. What number is displayed in place of the blurry text when you hold down the keys Q and ENTER?

Question 3:

Imagine that one of the mental arithmetic problems in part 2 is: **29 + 11 + 93 – 22**. What is the correct solution to this problem, for which you would receive a payoff of €0.10?

Imagine the following situation:

After a mental arithmetic problem in part 2, the following headline is displayed to you:
[screenshot of a screen with a headline]
The following two questions both refer to this situation. For each question, exactly one answer is correct.

Question 4:

What will happen if you click the button "Save this news article"?

◯ The news article will pop up immediately.
◯ The headline will disappear after about 10 seconds and the next mental arithmetic problem will start. Apart from that, nothing will happen.
◯ The headline will disappear immediately and the next mental arithmetic problem will start. Apart from that, nothing will happen.

◯ The headline will disappear after about 10 seconds and the next mental arithmetic problem will start. The news article will be redisplayed at the end of the study for me to read.

Question 5:

What will happen if you **don't** click the button "Save this news article"?

◯ The news article will pop up immediately.

◯ The headline will disappear after about 10 seconds and the next mental arithmetic problem will start. Apart from that, nothing will happen.

◯ The news article will be redisplayed at the end of the study for me to read.

Screen 4—Part 1.

[introductory sentence in the *Worry-Alleviating* condition:] Please watch the following video, which provides information about the current state of the preparations for vaccinations against COVID-19 in Germany.

[introductory sentence in the *Worry-Amplifying* condition:] Please watch the following video, which provides information about possible health damage caused by COVID-19.

Please click here in case the video doesn't load. [If subjects click on this text, they are instructed to watch the video on a linked external website.]

[embedded video player]

Please click "Continue" when you are done.

Screen 5—Question about the video.

[in the *Worry-Alleviating* condition:] According to the video, when could it be possible to carry out protective vaccinations against COVID-19 in Germany?

◯ only in a couple of years

◯ in summer 2021

◯ already in a few weeks

[in the *Worry-Amplifying* condition:] Which potential long-term health effects of COVID-19 are mentioned in the video?

◯ No long-term health effects are mentioned.

◯ Fatigue and shortness of breath can still occur after recovery.

◯ Even in individuals without pre-existing conditions, COVID-19 can cause muscle paralyses, strokes, and brain damage, among other things.

Screen 6—Part 1.

How would you rate your happiness at the moment? Please use a scale from 0 to 10, where 0 means "not happy at all" and 10 means "very happy".

How worried are you that COVID-19 could cause serious damage to your health or the health of your loved ones? Please use a scale from 0 to 10, where 0 means "not worried at all" and 10 means "very worried".

In the following part of the study, you will be able to earn an additional payoff by solving mental arithmetic problems. How many problems per block do you want to solve correctly? *Note:* In session 1 of the study, you answered [average] mental arithmetic problems per block correctly on average.

Screen 7—Part 2.

This part is divided into four blocks, each of which is 5 minutes long. During this time, you can once again solve mental arithmetic problems that consist of adding and subtracting four double-digit numbers. You will receive €0.10 for every correctly solved problem.
Sometimes, headlines of current news articles will be displayed to you for about 10 seconds in between two problems. If you want to read one of the displayed news articles, you can bookmark it for later by clicking the button "Save this news article". Bookmarked news articles will be redisplayed to you **at the end of the study**.

Note: If you don't click "Continue" before that, the first block will automatically start in **[countdown from 60]** seconds.

Screen 8.

Block 1 will start in a few seconds.

Screen 9.

[first 5-minute block of mental arithmetic problems]

Screen 10—Break.

[a relaxing image is displayed for 30 seconds]

Screen 11.

Block 2 will start in a few seconds.

Screen 12.

[second 5-minute block of mental arithmetic problems]

Screen 13—Part 2.

The first two blocks of mental arithmetic in part 2—block 1 and block 2—are over!

Do you think that you have solved more or fewer problems correctly on average in these two blocks than *yesterday in session 1 of the study*?
Note: Yesterday, you answered [average] problems per block correctly on average.

◯ In the two blocks now, I have on average solved fewer problems correctly than yesterday.
◯ In the two blocks now, I have on average solved about the same number of problems correctly as yesterday.
◯ In the two blocks now, I have on average solved more problems correctly than yesterday.

Please explain your answer in one to two sentences.

Screen 14.

Block 3 will start in a few seconds.

Screen 15.

[third 5-minute block of mental arithmetic problems, occasionally interrupted by the display of news headlines]

Screen 16—Break.

[a relaxing image is displayed for 30 seconds]

Screen 17.

Block 4 will start in a few seconds.

Screen 18.

[fourth 5-minute block of mental arithmetic problems, occasionally interrupted by the display of news headlines]

Screen 19—Part 2.

The last two blocks of mental arithmetic in part 2—block 3 and block 4—are over!

Do you think that you have solved more or fewer problems correctly on average in these two blocks than *yesterday in session 1 of the study*?

Note: Yesterday, you answered [average] problems per block correctly on average.

◯ In the two blocks now, I have on average solved fewer problems correctly than yesterday.
◯ In the two blocks now, I have on average solved about the same number of problems correctly as yesterday.
◯ In the two blocks now, I have on average solved more problems correctly than yesterday.

Please explain your answer in one to two sentences.

Screen 20—Questionnaire.

Finally, please answer a couple of questions that are relevant for the evaluation of the study.

Altogether, how satisfied are you with your performance in solving the arithmetic problems during **today's** session of the study? Please use a scale from 0 to 10, where 1 means "not satisfied at all" and 10 means "very satisfied".

At the beginning of today's session of the study, you have watched a video about the new coronavirus and subsequently answered a question about this video. The following three questions all refer to this video.

Question 1:

Do you still remember the main point of the video?

◯ yes
◯ no

Question 2:

Did the video contain information that was new to you?

◯ yes
◯ no

Question 3:

Was your focus on the mental arithmetic problems during today's session of the study impaired by thoughts about the content of the video?

◯ yes
◯ no

Do you actively follow the news about the new coronavirus (besides coming across news items casually)?

◯ yes
◯ no

Screen 21—Questionnaire (page 2).

Imagine that you could earn an additional payoff by working on a task. The task consists of alternately pressing the keys X and Y on your keyboard. To receive the payoff, you have to press a certain number of key combinations (pressing X and Y alternately once). Assume that you need one minute for 250 combinations. How many key combinations would you press for an additional payoff of **€2.00**?

How well do the following statements describe you as a person? Please use a scale from 0 to 10, where 0 means "does not describe me at all" and 10 means "describes me perfectly".

Statement 1:
"When someone does me a favor I am willing to return it."

Statement 2:
"If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so."

How worried are you about the impact of the COVID-19 pandemic on your own financial situation? Please use a scale from 0 to 10, where 0 means "not worried at all" and 10 means "very worried".

How worried are you about the impact of the COVID-19 pandemic on your job prospects? Please use a scale from 0 to 10, where 0 means "not worried at all" and 10 means "very worried".

Screen 22—Questionnaire (page 3).

What do you think? Out of 100 individuals in Germany who are currently not infected, how many will contract the new coronavirus within the next 6 months?

What do you think? Out of 100 individuals in Germany who contract the new coronavirus, how many will suffer from lasting health damage as a consequence of the disease?

What do you think? Out of 100 individuals in Germany who contract the new coronavirus, how many will die as a consequence of the disease?

What do you think? How many months will it take until a vaccine is available for the majority of the population in Germany?
Note: If you think that there won't be a vaccine, please fill in "-1".

Screen 23—Questionnaire (page 4).

[German version of the Penn State Worry Questionnaire (Glöckner-Rist and Rist, 2014)]

Screen 24—Feedback on the study.

If you want to, you can now provide feedback to the principal investigators of this study. Were the instructions clear? Did you experience difficulties in any part of the study? You can leave this field blank if you prefer not to give feedback.

Screen 25—Your payoff.

This online study is now over. Thank you for your participation!

Your payoff:
Your payoff is composed as follows:

- **€0.10** × **[count]** correctly solved arithmetic problems problems in total in sessions 1 and 2;

- **€4.00** for watching the video in session 2.

In total, you will therefore receive **[payoff]**.

You will receive your payoff by bank transfer. The bank transfer will be ordered no later than [bank transfer date], and the money should be available on your bank account a few days later.

In case of questions or problems regarding your payoff, please write an e-mail to lpcw-studie@uni-bonn.de.

Information about the coronavirus:

As part of the study, you watched a media report about the new coronavirus. Individual media reports may not cover all aspects of a given topic. Up-to-date, verified information about the coronavirus can be found on government websites. [If subjects click on this text, they are redirected to a government website with information about COVID-19.]

The news articles you bookmarked during the mental arithmetic problems:

To view bookmarked news articles, click on the respective headline. The articles will open in new tabs.

[all bookmarked news headlines with the corresponding links displayed one below the other]

Appendix 3.E Screenshots of the Experimental Interface

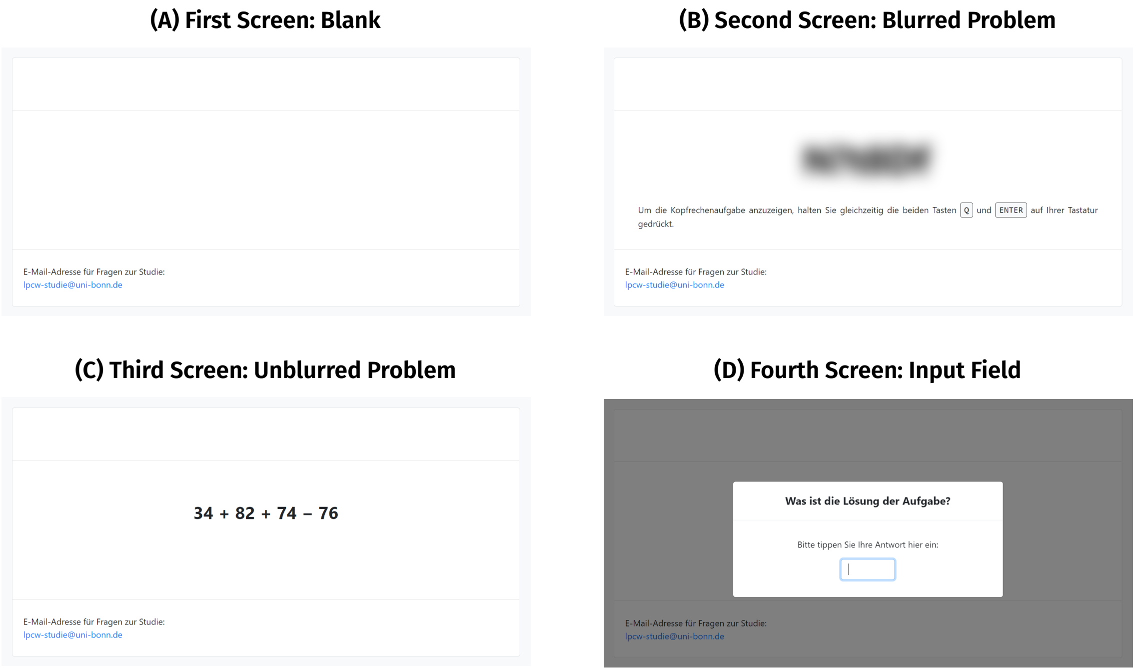

Figure 3.E.1. Sequence of Screens for Each Mental Arithmetic Problem of the Cognitive Task

Notes: Screenshots of the experimental interface for each screen in the sequence of events for a mental arithmetic problem, with text in German. The English translations for the text on the screenshots are “To display the mental arithmetic problem, simultaneously hold down the two keys Q and ENTER on your keyboard” for panel (B) and “What is the solution to the problem? Please type in your answer here:” for panel (D). In the bottom left corner of all screens, an e-mail address for questions about the study is provided.

References

Allen, Eric J., Patricia M. Dechow, Devin G. Pope, and George Wu. 2017. "Reference-Dependent Preferences: Evidence from Marathon Runners." *Management Science* 63 (6): 1657–72. DOI: 10.1287/mnsc.2015.2417. [120, 130]

Apenbrink, Christian. 2021. "The Cost of Worrying about an Epidemic: Ebola Concern and Cognitive Function in the US." ECONtribute Discussion Paper No. 120. Bonn/Cologne: University of Bonn and University of Cologne. [106, 110, 137]

Barber, Brad M., Wei Jiang, Adair Morse, Manju Puri, Heather Tookes, and Ingrid M. Werner. 2021. "What Explains Differences in Finance Research Productivity during the Pandemic?" *Journal of Finance* 76 (4): 1655–97. DOI: 10.1111/jofi.13028. [108, 109]

Bellet, Clement, Jan-Emmanuel De Neve, and George Ward. 2020. "Does Employee Happiness Have an Impact on Productivity?" Unpublished manuscript, Massachusetts Institute of Technology. December 17, 2020. URL: https://drive.google.com/open?id=1F2M8K8_9201VWPuhtQvGMotqYgH3q3CH. [138]

Bernstein, Shai, Timothy McQuade, and Richard R. Townsend. 2021. "Do Household Wealth Shocks Affect Productivity? Evidence from Innovative Workers During the Great Recession." *Journal of Finance* 76 (1): 57–111. DOI: 10.1111/jofi.12976. [105]

Blix, Ines, Marianne Skogbrott Birkeland, and Siri Thoresen. 2021. "Worry and Mental Health in the Covid-19 Pandemic: Vulnerability Factors in the General Norwegian Population." *BMC Public Health* 21: 928. DOI: 10.1186/s12889-021-10927-1. [140]

Bloom, Nicholas, Philip Bunn, Paul Mizen, Pawel Smietanka, and Gregory Thwaites. 2020. "The Impact of Covid-19 on Productivity." NBER Working Paper No. 28233. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w28233. [108]

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch. 2014. "hroot: Hamburg Registration and Organization Online Tool." *European Economic Review* 71: 117–20. DOI: 10.1016/j.euroecorev.2014.07.003. [117]

Bogliacino, Francesco, Cristiano Codagnone, Felipe Montealegre, Frans Folkvord, Camilo Gómez, Rafael Charris, Giovanni Liva, Francisco Lupiáñez-Villanueva, and Giuseppe A. Veltri. 2021. "Negative Shocks Predict Change in Cognitive Function and Preferences: Assessing the Negative Affect and Stress Hypothesis." *Scientific Reports* 11: 3546. DOI: 10.1038/s41598-021-83089-0. [137]

Borowiecki, Karol Jan. 2017. "How Are You, My Dearest Mozart? Well-Being and Creativity of Three Famous Composers Based on Their Letters." *Review of Economics and Statistics* 99 (4): 591–605. DOI: 10.1162/REST_a_00616. [138]

Brodeur, Abel, David Gray, Anik Islam, and Suraiya Bhuiyan. 2021. "A Literature Review of the Economics of COVID-19." *Journal of Economic Surveys* 35 (4): 1007–44. DOI: 10.1111/joes.12423. [108]

Brynjolfsson, Erik, John J. Horton, Adam Ozimek, Daniel Rock, Garima Sharma, and Hong-Yi TuYe. 2020. "COVID-19 and Remote Work: An Early Look at US Data." NBER Working Paper No. 27344. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w27344. [108]

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2021. "Opinions as Facts." Unpublished manuscript, University of Chicago. July 1, 2021. URL: https://home.uchicago.edu/bursztyn/OpinionsAsFacts.pdf. [109]

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang. 2016. "Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday." *American Economic Review* 106 (2): 260–84. DOI: 10.1257/aer.20140481. [110, 137]

Castriota, Stefano, Marco Delmastro, and Mirco Tonin. 2020. "National or Local? The Demand for News in Italy during COVID-19." IZA Discussion Paper No. 13805. Bonn: IZA – Institute of Labor Economics. URL: https://ftp.iza.org/dp13805.pdf. [109]

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. DOI: 10.1016/j.jbef.2015.12.001. [116]

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team. 2020. "How Did Covid-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data." NBER Working Paper No. 27431. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w27431. [108]

Clark, Damon, David Gill, Victoria Prowse, and Mark Rush. 2020. "Using Goals to Motivate College Students: Theory and Evidence from Field Experiments." *Review of Economics and Statistics* 102 (4): 648–63. DOI: 10.1162/rest_a_00864. [139]

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending." NBER Working Paper No. 27141. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w27141. [108]

Correia, Sergio, Stephan Luck, and Emil Verner. 2020. "Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu." Unpublished manuscript, Sloan School of Management, Massachusetts Institute of Technology. June 5, 2020. URL: https://www.emilverner.com/s/CorreiaLuckVerner_June11_2020.pdf. [109]

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman. 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." *Brookings Papers on Economic Activity* 51 (2): 35–82. DOI: 10.1353/eca.2020.0006. [108]

de Bruijn, Ernst-Jan, and Gerrit Antonides. 2022. "Poverty and Economic Decision Making: A Review of Scarcity Theory." *Theory and Decision* 92 (1): 5–37. DOI: 10.1007/s11238-021-09802-7. [137]

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, Amsterdam: North-Holland. Chapter 19, 723–68. DOI: 10.1016/B978-0-444-63685-0.00019-X. [109]

DeStefano, Diana, and Jo-Anne LeFevre. 2004. "The Role of Working Memory in Mental Arithmetic." *European Journal of Cognitive Psychology* 16 (3): 353–86. DOI: 10.1080/09541440244000328. [111]

Duong, Hai Long, Junhong Chu, and Dai Yao. Forthcoming. "Taxi Drivers' Response to Cancellations and No-Shows: New Evidence for Reference-Dependent Preference." *Management Science*, DOI: 10.1287/mnsc.2022.4349. [108, 120, 130]

Dupas, Pascaline, Jonathan Robinson, and Santiago Saavedra. 2020. "The Daily Grind: Cash Needs and Labor Supply." *Journal of Economic Behavior & Organization* 177: 399–414. DOI: 10.1016/j.jebo.2020.06.017. [120]

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2021. "The Macroeconomics of Epidemics." *Review of Financial Studies* 34 (11): 5149–87. DOI: 10.3386/w26882. [108]

Emanuel, Natalia, and Emma Harrington. 2021. "'Working' Remotely? Selection, Treatment, and Market Provision of Remote Work." Unpublished manuscript, Harvard University. April 9, 2021. URL: https://scholar.harvard.edu/files/eharrington/files/harrington_jmp_working_remotely.pdf. [109]

Eysenck, Michael W., Nazanin Derakshan, Rita Santos, and Manuel G. Calvo. 2007. "Anxiety and Cognitive Performance: Attentional Control Theory." *Emotion* 7 (2): 336–53. DOI: 10.1037/1528-3542.7.2.336. [106, 115, 120, 131]

Faia, Ester, Andreas Fuster, Vincenzo Pezone, and Basit Zafar. Forthcoming. "Biases in Information Selection and Processing: Survey Evidence from the Pandemic." *Review of Economics and Statistics*, DOI: 10.1162/rest_a_01187. [109]

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. 2016. "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences." IZA Discussion Paper No. 9674. Bonn: IZA – Institute of Labor Economics. URL: https://ftp.iza.org/dp9674.pdf. [116]

Fehr, Ernst, and Lorenz Goette. 2007. “Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment.” *American Economic Review* 97 (1): 298–317. DOI: 10.1257/aer.97.1.298. [108, 120]

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review* 98 (1): 394–425. DOI: 10.1257/aer.98.1.394. [105]

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth. 2021. “Work from Home & Productivity: Evidence from Personnel & Analytics Data on IT Professionals.” IZA Discussion Paper No. 14336. Bonn: IZA – Institute of Labor Economics. URL: https://ftp.iza.org/dp14336.pdf. [109]

Glöckner-Rist, Angelika, and Fred Rist. 2014. “Deutsche Version des Penn State Worry Questionnaire (PSWQ-d).” Zusammenstellung sozialwissenschaftlicher Items und Skalen (ZIS). DOI: 10.6102/zis219. [116, 167]

Goerg, Sebastian J., and Sebastian Kube. 2012. “Goals (Th)at Work: Goals, Monetary Incentives, and Workers' Performance.” Preprints of the Max Planck Institute for Research on Collective Goods Bonn 2012/19. Bonn: Max Planck Institute for Research on Collective Goods. DOI: https://homepage.coll.mpg.de/pdf_dat/2012_19online.pdf. [139]

Goette, Lorenz. 2021. “Reference Points and Effort.” Unpublished manuscript, University of Bonn. March 20, 2021. [120]

Goolsbee, Austan, and Chad Syverson. 2021. “Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020.” *Journal of Public Economics* 193: 104311. DOI: 10.1016/j.jpubeco.2020.104311. [108]

Hall, Robert E., and Charles I. Jones. 1999. “Why Do Some Countries Produce So Much More Output per Worker than Others?” *Quarterly Journal of Economics* 114 (1): 83–116. DOI: 10.1162/003355399555954. [105]

Heath, Chip, Richard P. Larrick, and George Wu. 1999. “Goals as Reference Points.” *Cognitive Psychology* 38 (1): 79–109. DOI: 10.1006/cogp.1998.0708. [120]

Kaufman, Bruce E. 1999. “Emotional Arousal as a Source of Bounded Rationality.” *Journal of Economic Behavior & Organization* 38 (2): 135–44. DOI: 10.1016/s0167-2681(99)00002-5. [106]

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2015. “Self-Control at Work.” *Journal of Political Economy* 123 (6): 1227–77. [139]

Kaur, Supreet, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach. 2021. “Do Financial Concerns Make Workers Less Productive?” NBER Working Paper No. 28338. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w28338. [107, 110, 129, 137]

Krekel, Christian, George Ward, and Jan-Emmanuel De Neve. 2019. “Employee Wellbeing, Productivity and Firm Performance.” CEP Discussion Paper No 1605. London: Centre for Economic Performance. URL: https://cep.lse.ac.uk/pubs/download/dp1605.pdf. [138]

Kruger, Samuel, Gonzalo Maturana, and Jordan Nickerson. 2020. “How Has COVID-19 Impacted Research Productivity in Economics and Finance?” Unpublished manuscript, Sloan School of Management, Massachusetts Institute of Technology. December 30, 2020. DOI: 10.2139/ssrn.3745226. [108]

Künn, Steffen, Christian Seel, and Dainis Zegners. 2022. “Cognitive Performance in Remote Work: Evidence from Professional Chess.” *Economic Journal* 132 (643): 1218–32. DOI: 10.1093/ej/ueab094. [109]

Lichand, Guilherme, and Anandi Mani. 2020. “Cognitive Droughts.” Department of Economics Working Paper Series No. 341. Zurich: University of Zurich. URL: https://www.econ.uzh.ch/static/wp/econwp341.pdf. [110, 115, 131, 137]

Lindqvist, Erik, and Roine Vestman. 2011. “The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment.” *American Economic Journal: Applied Economics* 3 (1): 101–28. DOI: 10.1257/app.3.1.101. [105]

Liu, Xiaoou, and Chen Zhu. 2014. “Will Knowing Diabetes Affect Labor Income? Evidence from a Natural Experiment.” *Economics Letters* 124 (1): 74–78. DOI: 10.1016/j.econlet.2014.04.019. [105]

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. “Poverty Impedes Cognitive Function.” *Science* 341 (6149): 976–80. DOI: 10.1126/science.1238041. [110, 137]

Markle, Alex, George Wu, Rebecca White, and Aaron Sackett. 2018. “Goals as Reference Points in Marathon Running: A Novel Test of Reference Dependence.” *Journal of Risk and Uncertainty* 56 (1): 19–50. DOI: 10.1007/s11166-018-9271-9. [120]

Mathews, Andrew. 1990. “Why Worry? The Cognitive Function of Anxiety.” *Behaviour Research and Therapy* 28 (6): 455–68. DOI: 10.1016/0005-7967(90)90132-3. [106, 120]

Meyer, Thomas J., Mark L. Miller, Richard L. Metzger, and Thomas D. Borkovec. 1990. “Development and Validation of the Penn State Worry Questionnaire.” *Behaviour Research and Therapy* 28 (6): 487–95. DOI: 10.1016/0005-7967(90)90135-6. [116]

Moran, Tim P. 2016. “Anxiety and Working Memory Capacity: A Meta-Analysis and Narrative Review.” *Psychological Bulletin* 142 (8): 831–64. DOI: 10.1037/bul0000051. [106, 120]

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means So Much*. New York, NY: Henry Holt and Company. [115, 121, 131, 137]

Neuhäuser, Markus, Anke Welz, and Graeme D. Ruxton. 2017. “Statistical Tests for the Comparison of Two Samples: The General Alternative.” *Communications in Statistics - Simulation and Computation* 46 (2): 903–9. DOI: 10.1080/03610918.2014.983651. [125]

Oswald, Andrew J., Eugenio Proto, and Daniel Sgroi. 2015. “Happiness and Productivity.” *Journal of Labor Economics* 33 (4): 789–822. DOI: 10.1086/681096. [106, 110, 111, 120, 138, 139]

Presse- und Informationsamt der Bundesregierung, Berlin. 2021. “Trend Questions Corona (Week 49/2020) [dataset].” Cologne: GESIS Data Archive. ZA7679 Data file Version 1.0.0 [distributor]. DOI: 10.4232/1.13704. [117]

Robinson, Oliver J., Katherine Vytal, Brian R. Cornwell, and Christian Grillon. 2013. “The Impact of Anxiety upon Cognition: Perspectives from Human Threat of Shock Studies.” *Frontiers in Human Neuroscience* 7: 203. DOI: 10.3389/fnhum.2013.00203. [106, 120]

Rosenbaum, Paul R. 2002. “Covariance Adjustment in Randomized Experiments and Observational Studies.” *Statistical Science* 17 (3): 286–327. DOI: 10.1214/ss/1042727942. [126]

Sacerdote, Bruce, Ranjan Sehgal, and Molly Cook. 2020. “Why Is All COVID-19 News Bad News?” NBER Working Paper No. 28110. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w28110. [109, 133]

Sari, Berna A., Ernst H. W. Koster, and Nazanin Derakshan. 2017. “The Effects of Active Worrying on Working Memory Capacity.” *Cognition and Emotion* 31 (5): 995–1003. DOI: 10.1080/02699931.2016.1170668. [106, 120]

Shah, Anuj K., Sendhil Mullainathan, and Eldar Shafir. 2012. “Some Consequences of Having Too Little.” *Science* 338 (6107): 682–85. DOI: 10.1126/science.1229223. [137]

Sheridan, Adam, Asger Lau Andersen, Emil Toft Hansen, and Niels Johannesen. 2020. “Social Distancing Laws Cause Only Small Losses of Economic Activity during the COVID-19 Pandemic in Scandinavia.” *Proceedings of the National Academy of Sciences* 117 (34): 20468–73. DOI: 10.1073/pnas.2010068117. [108]

Simonov, Andrey, Szymon Sacher, Jean-Pierre Dubé, and Shirsho Biswas. 2022. “Frontiers: The Persuasive Effect of Fox News: Noncompliance with Social Distancing during the COVID-19 Pandemic.” *Marketing Science* 42 (2): 230–42. DOI: 10.1287/mksc.2021.1328. [109]

Snowberg, Erik, and Leeat Yariv. 2021. “Testing the Waters: Behavior across Participant Pools.” *American Economic Review* 111 (2): 687–719. DOI: 10.1257/AER.20181065. [117]

van den Berg, Gerard J., Petter Lundborg, and Johan Vikström. 2017. “The Economics of Grief.” *Economic Journal* 127 (604): 1794–832. DOI: 10.1111/ecoj.12399. [105]

Zeckhauser, Richard. 1996. “The Economics of Catastrophes.” *Journal of Risk and Uncertainty* 12 (2–3): 113–40. DOI: 10.1007/BF00055789. [109]

Zizzo, Daniel John. 2010. “Experimenter Demand Effects in Economic Experiments.” *Experimental Economics* 13 (1): 75–98. DOI: 10.1007/s10683-009-9230-z. [134]