

On Avoidance and Neglect as Barriers to Informed Decision-Making

Three Essays in Behavioral Economics

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
MARKUS PETER FELS
aus Berlin

Bonn 2014

Dekan: Prof. Dr. Klaus Sandmann
Erstreferent: Prof. Dr. Paul Heidhues
Zweitreferent: Prof. Dr. Daniel Krähler

Tag der mündlichen Prüfung: 23.07.2014

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn
(http://hss.ulb.uni-bonn.de/diss_online) elektronisch publiziert.

Acknowledgements

I am grateful for the support I received from many people during the time I worked on this thesis.

First, I want to thank Paul Heidhues for arousing my interest in Behavioral Economics and for his great supervision. I benefited a lot from his advice. I also want to thank Daniel Krähmer for being my second supervisor and for providing many helpful comments and suggestions to improve my work. I am thankful to Dezsö Szalay for joining my dissertation committee. I am grateful to Matthias Kräkel for helpful feedback on my work.

I want to thank Carsten Dahremöller for an extraordinary collaboration on a project that is the basis of Chapter II. The joint work was inspiring and fun. I am grateful for our many discussions, his enthusiasm, and his patience.

The Bonn Graduate School of Economics (BGSE) provided an ideal environment for my Ph.D. studies as well as financial support. I am grateful to all the people that help to create and maintain this great environment for learning and research, in particular Urs Schweizer and Silke Kinzig.

I owe thanks to my fellow BGSE students, in particular Sina Litterscheid who shared an office with me for many years. I experienced a remarkable atmosphere of cooperation and friendship, of mutual assistance and motivation during my time at the graduate school.

I thank Philipp Reiss for his advice and motivation during the final period of my dissertation.

Finally, I want to thank my family and friends for their love and support.

Contents

I. On the Value of Information: Why People Reject Medical Tests	5
1. Introduction	6
2. The Model	9
3. The Value of Information	22
4. Screening and the Test as Gate-Keeper	28
5. Comparative Statics	31
6. Conclusion	34
II. Product Lines, Product Design, and Limited Attention	37
1. Introduction	38
2. The Role of Attention	41
3. Optimal Product Design of a Monopolist	48
4. Optimal Product Design with Introduction of a Bait Good	51
5. Discussion	54
6. Conclusion	56
III. Limited Attention and the Demand for Health Insurance	57
1. Introduction	58
2. The Problem of Buying Insurance	61
3. Profitable Undercutting in an Insurance Market	69
4. Dominated Choices	73
5. Specific vs. Comprehensive Insurance: The Benefits of Isolating Risks	75
6. Insurance Provision in Market Equilibrium	77
7. Conclusion	80
A. Appendices	82
1. Appendix to Chapter I	82
2. Appendix to Chapter II	92
3. Appendix to Chapter III	102

List of Figures

I.1. Difference between Posterior Beliefs when Information is Noninstrumental (default action NT)	19
I.2. Difference between Posterior Beliefs when Information is Instrumental .	21
A.1. Comparing F and G versus comparing m_F and m_G	84

Introduction

Classic economic argument ties the value of information to its potential to change future decisions. A decision-maker reaps no benefit from being informed better if this does not entail a new course of action. On the other hand, a decision-maker profits - in expectation - from being informed better whenever the new information makes him change his mind about some subsequent action choice. In this view, information may not hurt a decision-maker, or, as Radner and Stiglitz put it, “information is harmless”¹. The sole reason for people to remain uninformed is the limited or costly availability of information. In consequence, whenever undesirable outcomes are or have been linked to a lack of sufficiently informed decision-making one of the most popular policy recommendations is to provide free information.

This work collects three essays on two barriers to informed decision-making. In particular, it investigates potential reasons for making choices that are un-, or at least less informed than possible, despite the free availability of information.

Information avoidance is the first barrier that is discussed. When being offered free information people may be unwilling to acquire information and prefer to make decisions in ignorance of it. Such behavior suggests costs to information acquisition beyond the costs of producing or researching information. This work considers psychological costs associated with information acquisition, namely the emotional reaction that is anticipated to be triggered by information reception.

The second barrier to informed decision-making that is investigated is the neglect of available information. While a person may be willing to base his decisions on as much information as possible that person may be simply unable to do so. Given the availability of information such inability must then be the result of an inaptitude to process the information. This work considers limitations in cognitive resources to process large amounts of information as the source of such inaptitude. A full incorporation of all available information relevant to a decision may be well beyond the cognitive means of a decision-maker. Alternatively, the cognitive costs of an incorporation of all available

¹Radner and Stiglitz (1984), p.33.

information may fall short of the incremental benefit from making a fully as opposed to a slightly less informed decision. As a result the decision-maker neglects some part of the available information. This thesis thus highlights two different reasons for making less-informed judgments and investigates some of their implications.

Chapter 1 investigates active avoidance of information as a source of uninformed decision-making. Next to providing input for subsequent decisions information may trigger undesirable emotional responses. An optimal decision on information acquisition then weighs benefits from better decision-making against the costs of an inferior emotional well-being.

The essay investigates the decision to undergo a medical test. Although such tests convey important information, a significant number of people are found to turn down an opportunity to test for free. The essay employs a model proposed by Kőszegi and Rabin (2009) to determine news utility, a concept suitable to depict the emotional response to information. This allows deriving the value of information as a function of the decision-maker's prior belief and the information's precision. Information gains value from enhancing decision-making, as proposed by the classic view, but also from alleviating adverse emotional responses in the future. It loses value by potentially triggering adverse emotional responses upon reception.

Choosing parameter values in a way to approximate the setting of a screening test, a condition is derived under which rejection of testing is optimal. This condition is satisfied when the decision-maker is offered to be tested for a severe disease for which treatment is beneficial, yet ineffective in curing the disease. This best describes situations in which care, but no cure is available. Comparative statics reveal the timing of tests to be another major determinant of test uptake. Both theoretical predictions are well in line with stated reasons to decline testing.

The chapter derives the value of information for a person with reference-dependent preferences in the context of medical decision-making. It discusses the interdependency between the value of information from a decision-making and an emotional point of view. Finally, it derives testable predictions on the desirability of information and test uptake.

The following two chapters investigate the (partial) neglect of information that results from the inaptitude of people to process large amounts of information. Here, uninformed or less-informed decision-making is not the result of active avoidance. It is the lack of the cognitive capacity to process more information that results in decisions that are less than fully informed.

Chapter 2, based on cooperation with Carsten Dahremöller, proposes a model that depicts how a decision-maker simplifies a multi-dimensional choice problem in order to be able to make such a complex decision. The model posits that the decision-maker focuses on those dimensions in which differences are deemed most important, and those in which the available alternatives differ the most. The resulting attention allocation reflects scarcity of attention through neglect. Further, it reflects an underlying optimality in the allocation of attentional resources as focus is drawn towards those dimensions that are most important.

The chapter further investigates how a monopolist optimally designs her products when facing customers with such limited attention. The resulting design features simplicity. A customer may value additional qualities of a product, yet if he does not incorporate those in his purchase decision, i.e. if they do not translate into a higher willingness-to-pay for the product, it is not optimal to include those qualities into the product design. Further it is shown that the monopolist tends to suffer from the limited attention of her customers while the customers are generally better off. This is in strong contrast to the previous literature that highlights the exploitability of inattentive customers.

It is further argued that the monopolist may, under quite general conditions, profitably employ bait goods. These products are solely designed to increase the customers' attention and to draw attention to more profitable attributes while themselves being highly desirable, yet too expensive to purchase. Since these bait goods tend to be high-end, state-of-the-art products the monopolist increases the willingness-to-pay for the more moderate primary good via a compromise effect. The model thus offers an alternative to the existing approaches explaining the effect based on reference-dependent preferences or contextual inference. Finally, it is argued that the optimal attention manipulation does not generally draw most attention to the most profitable attributes. The second chapter proposes a model of limited attention based on the assumption that a decision-maker needs to simplify a complex decision in order to make it. It highlights the shortcomings but also the benefits of simplified decision-making. In an application to the question of optimal product design, it derives the implications of such limited attention in a market setting and offers a novel explanation for the compromise effect.

The final chapter applies the model of Chapter 2 to the setting of health plan choice. This choice situation is often associated with high complexity since plans may differ on a large number of potentially relevant aspects. The most distinguishing feature from the previous chapter is the focus on a setting in which choices are made under uncertainty. In fact, most of the complexity involved in health plan choice is based on this

uncertainty. Further, the chapter seeks to show how the model may explain empirical evidence from the domain of health plan choice.

The model of limited attention predicts an undervaluation of insurance by inattentive customers. The advantages of plan purchase are distributed across many dimensions while the disadvantage of plan purchase is concentrated in a single dimension: the premium. This leads customers to focus on the premium dimension while not taking into account all of the benefits of being insured. This can lead to insurance not being purchased despite it being individually optimal. By a similar logic, the model predicts low quality, low premium plans to be suboptimally attractive to inattentive customers. Both predictions are well in line with empirical evidence from Medicare Part D plan choice. Another observation that has been made in the domain of health plan choice and that is hard to reconcile within classic models are dominated choices. The model predicts such behavior as customers may happen to neglect exactly those dimensions in which the domination occurs. Finally, the model predicts submodularity in the willingness-to-pay for insurance, a prediction that conforms to experimental evidence on insurance choice.

The last chapter shows how the model of limited attention may explain diverse empirical patterns that have been observed in the choice of health plans. Furthermore, it is a first attempt to discuss the complexity inherent in decisions under uncertainty and its implications. Finally, it is a showcase example for a setting in which uninformed or badly-informed decisions can be observed despite, or, as the model suggests, exactly because of an abundance of available information.

The essays collected here discuss two barriers to informed decision-making despite free availability of information: avoidance and neglect. Both the reasons for and several implications of such behavior are discussed. Although the free provision of information often seems a promising intervention to further informed decision-making, the essays collected here may warn against it being considered a panacea.

I. On the Value of Information: Why People Reject Medical Tests

In this chapter a model of reference-dependent preferences proposed by Kőszegi and Rabin (2009) is used to derive the value of information when a decision-maker is loss averse over changes in beliefs. This allows to model the anticipation of potential disappointment upon receiving bad news. It is shown that this emotional impact changes if information is instrumental, i.e. if it is affecting the decision about a subsequent action. The questions whether information is desirable from a decision-making or from an emotional point of view can thus not be separated. The model is applied to a patient's choice problem to undergo medical screening. The availability of effective cure and the timing of testing are predicted to be significant determinants of test uptake. This is in line with empirical research concerning patients' motives to decline testing.

1. INTRODUCTION

Medical diagnosis is an important part of health care provision. Tests are conducted to guide medical decision-making, particularly to identify the very need of a medical intervention. The actual use of medical information, or the lack of it, has created a couple of puzzles though. One of them is the unwillingness of some patients to take a medical test. Despite extensive campaigns to raise awareness people are reluctant to take up screening tests for breast or colon cancer. Similarly, most people acknowledge the benefits of testing for genetic predispositions for diseases such as Huntington's disease or breast cancer, yet uptake rates fall far below expectations. Finally, refusal of HIV tests remain a concern despite awareness campaigns, the offer of free and anonymous tests, and even the general acknowledgment of the benefits of these tests on the side of both physicians and patients. These observations are not easy to reconcile with the predictions of standard decision theory concerning the value of information. As a result, psychological barriers to testing have drawn increasing attention as possible explanations.¹ Psychological motives are among the most frequently reported reasons to decline testing. Reasons such as "fear of knowing one's status", "fear of a positive test result", "concern about the ability to cope with a positive result", and "emotional reactions" are commonly cited. Some studies suggest that it is the unavailability of effective treatment that drives the fear of being tested positive.² The anticipation of the emotional impact of information thus seems to be an important factor determining attitudes toward information. This work seeks to investigate this particular nuance of attitudes toward information and its interaction with the desirability of information in terms of improved decision-making.

The model we employ was introduced by Kőszegi and Rabin (2009), henceforth KR. It suggests that people derive (dis)utility from changes in beliefs about future outcomes. We interpret this utility from a change in belief as the emotional reaction to information. This enables us to model different incentives for information acquisition. On the one

¹See e.g. Neumann, Hammitt, Mueller, Fillit, Hill, Tetteh, and Kosik (2001), Lerman, Seay, Balshem, and Audrain (1995), and Geer, Ropka, Cohn, Jones, and Miesfeldt (2001) on barriers to testing for genetic predispositions for Alzheimer's Disease, Huntington's Disease and various forms of cancer. See Deblonde, Koker, Hamers, Fontaine, Luchters, and Temmerman (2010) for a review of studies on HIV testing in Europe. See Weiser, Heisler, Leiter, de Korte, Tlou, DeMonner, Phaladze, Bangsberg, and Iacopino (2001) as an example of a study on HIV testing in Africa (Botswana) and Zapka, Stoddard, Zorn, McCusker, and Mayer (1991) as an example of a study on HIV testing in the North America (USA). These studies are selected as they report reasons given by subjects for obtaining or rejecting an HIV test.

²See e.g. Neumann, Hammitt, Mueller, Fillit, Hill, Tetteh, and Kosik (2001), and Zapka, Stoddard, Zorn, McCusker, and Mayer (1991).

side, information is an input to subsequent decision-making, here with regard to a treatment choice. On the other hand, information may trigger an unfavorable emotional response. We derive the value of information as a composite of its value in terms of improved decision-making and its value in terms of emotional self-management. We find that these two components are interdependent, thus instrumentality and emotionality of information affect each other. In addition, the value function enables us to make predictions when to expect test refusal. It turns out that treatment effectiveness, in addition to treatment efficiency, plays a major role for the uptake of screening tests. This is because, as suggested by survey responses, the expected emotional response to the information provided by a screening test is particularly severe when there is no effective treatment available. Finally, the timing of testing turns out to be decisive for the desirability of testing as it affects the intensity of the emotional response.

This work joins the growing literature on reference-dependent preferences. The idea of outcomes being evaluated relative to a reference point has been prominent since Kahneman and Tversky (1979). Kőszegi and Rabin (2006) and Kőszegi and Rabin (2007) suggest to endogenize the reference point as being previously held expectations.³ While empirical research has found support for this hypothesis (Post, van den Assem, Baltussen, and Thaler (2008), Abeler, Falk, Goette, and Huffman (2011), Crawford and Meng (2011), Gill and Prowse (2012)), theoretical work has focused on the implications of such preferences (see e.g. Heidhues and Kőszegi (2008), Heidhues and Kőszegi (2010) on pricing strategies, and Herweg, Müller, and Weinschenk (2010) on optimal contracts). Kőszegi and Rabin (2009) further extend the model to account for utility being derived from changes in expectations. It is this extension that allows us to model emotions that result from changes in beliefs, in particular from the reception of information. KR use the model themselves to investigate information preferences. In contrast to this work they concentrate on noninstrumental information, i.e. information that is not affecting subsequent decisions. When they allow information to affect decision-making in a consumption-and-savings model they concentrate on the question how the subsequent action choice is affected, but leave out the question concerning the desirability of information itself. Karlsson, Loewenstein, and Seppi (2009) propose a model of selective attention in which an investor can influence the speed of adjustment

³In contrast to models of disappointment in which the reference point is the certainty equivalent of the lottery, such as Bell (1985) and Loomes and Sugden (1986) among others, the reference “point” is characterized by the whole lottery that represents the decision-maker’s expectations. An outcome’s evaluation can then lead to mixed feelings in the sense that an intermediate outcome compares worse against a better but better against a worse counterfactual. See Kőszegi and Rabin (2007) and Herweg, Müller, and Weinschenk (2010) for more elaborate discussions of this distinction from disappointment theories à la Bell (1985) and Loomes and Sugden (1986).

of his reference point by paying more or less attention to information. They confine analysis to noninstrumental information suggesting that if the information serves as an input to subsequent decision-making it gains an additional option value. We show in this chapter that this suggestion neglects the interdependence between an information's value in terms of instrumentality and emotionality. Close to our work Panidi (2008) analyses patients' desire to visit the doctor using a reference-dependent preferences framework. Similar to Karlsson, Loewenstein, and Seppi (2009) she neglects how the existence of a subsequent choice (here the possibility to treat), to which the information serves as input, changes the emotional impact of the information. Matthey (2008) proposes a different model of utility being derived from changes in beliefs, which she calls "adjustment utility". Apart from the observation that this adjustment utility may induce a distaste for positive but false information, Matthey (2008) focuses on the implications of such preferences for subsequent action choices but neglects their impact on information choice which is the main topic of this chapter.⁴

A different strand of literature that investigates the role of emotions in people's demand for information uses the concept of *anticipatory utility* (Kőszegi (2003), Caplin and Eliaz (2003), Caplin and Leahy (2004), Kőszegi (2006), Barigozzi and Levaggi (2008), Barigozzi and Levaggi (2010)). It suggests that individuals derive utility from holding specific beliefs. The psychological motive that affects informational choices is thus the maintenance of a positive - though potentially illusory - belief and not the avoidance of psychological distress (disappointment) resulting from bad news. The difference amounts to modeling different emotions as potential reasons for information avoidance. Anticipatory utility models *anticipatory feelings* such as anxiety, hope, or suspense. In contrast, this work seeks to model *anticipated feelings*, in particular the anticipation of shock, disappointment, or relief as a response to information reception. The distinction between anticipatory feelings and anticipated feelings is discussed by Loewenstein, Weber, Hsee, and Welch (2007). In addition to a different approach as to which emotions are modeled, this chapter highlights a trade-off that has been neglected in the literature on anticipatory utility. When deciding whether to acquire information, the decision-maker does not only face a trade-off between managing his emotions today and making better decisions (that pay off tomorrow). As we consider a situation in which individuals learn the truth eventually, information acquisition also has implications for emotional states or reactions tomorrow. Thus, a decision-maker faces another trade-off between managing today's emotions versus managing

⁴Interestingly though, Matthey (2008) highlights the propensity of such preferences to induce a psychological cost associated to deviations from previously made plans. Although these are not within the focus of this work such penalties from deviations can be observed at different instances in the analytical part of this chapter.

tomorrow's emotions. A direct implication of neglecting this trade-off is to ascribe negative value to information that does not influence decision-making.⁵ Finally, the literature on anticipatory utility models information aversion as the primary incentive for information refusal.⁶ As such, it can only identify the degree of information aversion and the decision-making value of information (i.e. the cost of an inefficient action), as reasons for information refusal. The first, in particular, is hardly a variable on which to base advice for practitioners who seek to diminish test refusal. In contrast, the model presented in this chapter explicitly derives and discusses which factors drive emotional reactions to information, and thus allows some insight into potential policy variables, such as the speed of tests. Admittedly, this comes at the expense of a more complex derivation of the value of information. Yet, we regard the additional insight concerning the factors that drive emotional responses and ultimately test refusal to be worth that cost as they suggest policy variables. Since we regard the literature on anticipatory utility to be very close to the idea discussed here, we will underline differing predictions at several instances.

This work thus seeks to complement the existing literature by investigating individuals' inclination to take a medical test if the information conveyed by the test both serves as an input to subsequent decision-making and triggers an emotional response. It highlights the interdependence between these two consequences of information choice and makes predictions concerning the determinants of test refusal.

The chapter is organized as follows. Section 1 introduces the model. In section 2, we derive the value of information as a function of prior beliefs and the testing technology. At the end of this section, we apply the model to an analytically simple case, the value of perfect information. Section 3 applies the model to the case of a screening test, and derives a condition under which rejection of a screening test is optimal. Section 4 shows comparative statics in order to illustrate how different parameters, such as treatment benefits and costs, severeness of disease, and the timing of the testing affect test refusal and, more generally, the value of information. Section 5 concludes.

2. THE MODEL

A decision-maker faces a problem spanning over three periods $t = 0, 1, 2$. In period 0, he has the possibility to test for a disease. Let $b \in \{i, n\}$ denote the information choice where i denotes the decision to test, n denotes the decision not to test. If he decides in favor of the test he will receive its result in the following period, period 1. Regardless of his information choice he has to make an action choice $a \in \{NT, T\}$ in period 1, where

⁵In contrast, proposition I.1 describes the possibility of a positive value of such information.

⁶Technically, information aversion means that the utility function over beliefs is concave.

NT denotes “no treatment” and T denotes “treatment”. He will make this treatment decision based on his period 1-belief about his health status. If he decided in favor of the test this belief will contain the information conveyed by the test. If he declined to take the test he has to choose an action without further information. In the final period (period 2), health outcomes realize based on his action choice and his health status.

There are two states of nature $\theta \in \{\theta_h, \theta_s\}$, meaning healthy (θ_h) and sick (θ_s). Let $p_0 \in (0, 1)$ be the subjective (prior) probability an individual assigns to being sick when making his information choice in period 0. The individual decides whether to obtain a signal s , a medical test, that conveys one of two possible messages $\{s^-, s^+\}$. The signal $s^-(s^+)$ is conclusive towards state θ_h (θ_s), i.e. s^- denotes a negative, s^+ a positive test result. If the patient chooses to be tested he expects to receive a negative (positive) test result with probability $q^-(q^+)$. Let $p_1 \in \{p^-, p_0, p^+\}$ denote the (posterior) probability an individual assigns to being sick in period 1, i.e. after potential information reception, where $p^-(p^+)$ is the posterior probability the individual assigns to being sick after having received a negative (positive) test result. Furthermore, let $\epsilon^-(\epsilon^+)$ denote the false negative (false positive) rate of the test, i.e. the probability of receiving message $s^-(s^+)$ given state θ_s (θ_h). It is assumed that these error rates are objective characteristics of the signal, i.e. they are known statistics, with $\epsilon^- + \epsilon^+ \leq 1$. Thus, the posterior (state) probabilities p^-, p^+ and the probabilities with which the decision-maker expects to receive each test result, denoted by q^-, q^+ , can be calculated via Bayes’ rule given the characteristics of the signal (ϵ^-, ϵ^+) and the subjective belief p_0 of an individual.

$$\begin{aligned}
 p^- &= \frac{\epsilon^- p_0}{q^-}, & p^+ &= \frac{(1 - \epsilon^-) p_0}{q^+} \\
 q^- &= (1 - \epsilon^+)(1 - p_0) + \epsilon^- p_0, & q^+ &= \epsilon^+(1 - p_0) + (1 - \epsilon^-) p_0
 \end{aligned}$$

Conditional on the state obtaining, the two actions (no treatment, treatment) lead to different levels of *material utility* $m \in \mathbb{R}$ in period 2, henceforth called payoffs or material outcomes. Let these be

	θ_h	θ_s
NT	A	B
T	C	D

Assume an ordering of payoffs $A > C > D > B$ and define the following differences between payoffs.

First, the differences across actions given a state are denoted by

$$\Delta_h = m(\text{no treatment, healthy}) - m(\text{treatment, healthy}) = A - C > 0$$

$$\Delta_s = m(\text{treatment, sick}) - m(\text{no treatment, sick}) = D - B > 0.$$

Δ_h denotes the net benefit of not being treated (or the net cost of being treated) to a healthy patient. For simplicity, we assume that treatment yields no benefit to a healthy individual. This allows Δ_h to be interpreted as the costs of treatment such as potential side effects and/or unpleasantness of the treatment procedure. Δ_s denotes the net benefit of treatment to a sick individual. Both Δ_h and Δ_s are strictly positive, so the optimal action differs across states. Assuming treatment costs to be independent across states, the sum ($\Delta_h + \Delta_s$) can be interpreted as the (gross) benefit of treatment to a sick.

Second, define the differences in utility across states given an action by

$$r_{NT} = m(\text{no treatment, healthy}) - m(\text{no treatment, sick}) = A - B > 0$$

$$r_T = m(\text{treatment, healthy}) - m(\text{treatment, sick}) = C - D \geq 0.$$

These will be important for determining the psychological evaluation of an outcome. The difference in well-being of an untreated healthy and an untreated sick individual (r_{NT}) can be interpreted as a measure of the severeness of the disease. The difference in well-being of treated healthy and a treated sick individual is denoted by r_T . It is a function of the *effectiveness* of treatment. If $r_T = 0$, the treatment constitutes a perfect cure as a treated sick is as well off as a treated healthy individual.⁷ The larger r_T the less effective the treatment.

In classic decision theory the optimal action choice as well as the optimal information choice, and thus the value of information, are a function of Δ_h , Δ_s , and the beliefs p_0 and p_1 of the decision-maker (DM). In contrast, if the DM exhibits reference-dependent preferences (RDP), they depend on r_{NT} and r_T in addition. We will now discuss the preferences of the decision-maker.

Incentives for Decision-Making

In each period, the forward-looking decision-maker has an objective function encompassing the utility in the current and all future periods.

$$U^t = \sum_{\tau=t}^2 u_{\tau}, \quad t = 0, 1, 2$$

⁷Note that a perfect cure is not necessarily costless.

There are two sources of utility. First, material utility is derived from state-contingent consequences of the action taken as was described above. Second, there is gain-loss utility derived from changes in belief regarding material utility. This means that news regarding one's future well-being affect well-being today, or, in the current context, getting a positive or negative test result affects a person's utility in the period these results are obtained. This gain-loss utility will be interpreted as elation or disappointment upon reception of good or bad news.

Changes in the belief about the future health outcome m can occur in all three periods. They can be the result of the reception of new information: in period 1 when a test result is received and in period 2 when the true state of nature is revealed. They can also result from a deviation of actual from planned choice: in period 0 when the information choice is made and in period 1 when the treatment choice is made.

The utility function will be modeled in a way proposed by Kőszegi and Rabin (2009)⁸, where the utility derived from changes in beliefs about future outcomes is called *prospective gain-loss utility* (PGLU), and the utility derived from changes in beliefs about current outcomes is called *contemporaneous gain-loss utility* (CGLU). An individual's utility in period t is an additively-separable function of material utility and gain-loss utility obtaining in this period

$$\begin{aligned} u_0 &= \gamma_0 v(F_0, F_{-1}), \\ u_1 &= \gamma_1 v(F_1, F_0), \\ u_2 &= m + v(F_2, F_1), \end{aligned}$$

where $m \in \mathbb{R}$ is the level of material utility occurring in period 2, $v(\cdot)$ is the level of gain-loss utility resulting from a change in belief about m , $F_t : \mathbb{R} \rightarrow [0, 1]$, $t = 0, 1, 2$ is the belief the patient holds at the end of period t concerning the level of material utility m in period 2, F_{-1} is the belief regarding material utility held immediately prior to information choice in period 0, and $\gamma_0, \gamma_1 \in [0, 1]$ are coefficients weighting the relative impact of prospective gain-loss utility compared to contemporaneous gain-loss utility. Following a suggestion by KR, we assume $0 \leq \gamma_0 \leq \gamma_1 \leq 1$, indicating a decline in the impact of changes in belief on well-being the larger the distance between the time the change in belief occurs and the time the material utility to which the belief refers is realized.

Following KR, the gain-loss utility from a change in belief is a function of how the new belief compares against the old belief. Formally, let $F : \mathbb{R} \rightarrow [0, 1]$ be a cumulative

⁸For simplicity we only consider a single-dimensional outcome in period 2, say utility from health status. While health status is clearly a multi-dimensional objective it will be treated as single-dimensional here for the sake of simpler exposition.

density function over material utility levels and define the function $v : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$, where \mathcal{F} is the set of all possible cumulative density functions $F(m)$ over material utility m , by

$$v(F_t, F_{t-1}) = \eta \int_{-\infty}^{\infty} \mu [F_{t-1}(m) - F_t(m)] dm$$

where η is a coefficient weighting gain-loss utility in relation to material utility. Analysis is simplified by assuming a two-piece linear representation of μ :

$$\begin{aligned} \mu(z) &= z \text{ if } z \geq 0 \\ \mu(z) &= \lambda z \text{ if } z < 0 \end{aligned}$$

with $\lambda > 1$ measuring the degree of loss aversion.⁹

Planning Behavior

Contemplating the decision whether to take the test or not, a forward-looking DM with above preferences realizes that this decision can affect his utility through three channels. First, the information conveyed by the test may affect his decision whether to treat, thereby affecting material and gain-loss utility in the final period. Second, even if the treatment choice remains unaffected, gain-loss utility in the final period will be affected. As this gain-loss utility is derived from the resolution of the remaining uncertainty, the information choice influences this source of utility by changing how much uncertainty there is prevailing until the final period. Third, as the information conveyed by the test changes his belief about his health he experiences gain-loss utility from the reception of information. In other words, he experiences an emotional response upon information reception.

In addition to contemplating the consequences of his behavior, the DM needs to keep in mind two things when laying out his plan how to behave in both instances of choice. First, he cannot rationally plan to make a choice that is not optimal at the time of choice. Second, once he has set up a plan, he will form beliefs about his future health status. Any future deviation from the plan will then result in a change in this belief producing gain-loss utility at the time of the deviation.

KR have proposed the concepts of personal equilibrium and preferred personal equilibrium to accommodate the restrictions just described. Formally, denote by d_t a “plan” for behavior starting in period t , i.e. a state-contingent strategy regarding all possible

⁹The exact approach of KR is define a function m_F that attributes levels of material utility to each percentile of a belief F and then compare the functions associated with two beliefs percentile by percentile. Under the assumption of two-piece linearity of μ both approaches yield the same gain-loss utility while the approach employed here spares the detour of deriving m_F . A proof of equivalence is in the Appendix.

decisions that might occur in period t and thereafter. Let D_t be the set of feasible plans.

Definition I.1. ¹⁰ Define the sets $\{D_t^*\}_{t=0}^T$ in the following backward-recursive way. A plan $d_t \in D_t$ is in D_t^* if, given the expectations generated by d_t , in any contingency, (i) it prescribes a continuation plan in D_{t+1}^* that maximizes the expectation of U^t , and (ii) it prescribes an action in period t that maximizes the expectation of U^t , assuming that future plans are made according to (i). A plan $d_t \in D_t$ is a personal equilibrium (PE) if $d_t \in D_t^*$. A plan $d_1 \in D_1$ is a preferred personal equilibrium (PPE) if $d_1 \in D_1^*$ and it maximizes the expectation of U^1 among plans in D_1^* .

Adapted to our setting this means that the plan d_1 prescribes a state-contingent treatment decision $a \in \{NT, T\}$ in period 1. This plan being a personal equilibrium requires that for any information set at which the DM finds himself in period 1 it must be optimal to follow the prescription a the plan makes for this information set given the expectations the DM with plan d_1 has at this information set. To give an example: a plan can only prescribe to seek treatment after a positive test result if, given the belief jointly determined by the information of the positive test result and the plan of the DM to treat in this instance, the DM finds it optimal to choose treatment. One step earlier, a plan d_0 comprises a planned information choice and a vector of planned treatment choices, $d_0 = (b, \vec{a})$. The vector of planned treatment choice \vec{a} comprises the planned treatment choice a for each possible information set in period 1. To give two examples: a plan d_0 may prescribe to test, and seek treatment upon a positive and abstain from treatment upon a negative test result; or it may prescribe not to test and abstain from treatment. Again, for d_0 to be a personal equilibrium it must be optimal to follow this plan given its anticipation. Note that more than one plan can constitute a personal equilibrium. In this case we assume that the DM chooses his preferred personal equilibrium, that is the plan that maximizes the expected utility from its implementation.

With these restrictions in mind, we now consider the three channels through which the information choice may affect utility. First, we will investigate which treatment choices can be part of personal equilibria. This is important as an information's propensity to influence the action choice constitute an important part of its value.

The Treatment Decision

In period 1, the DM decides whether to seek treatment. He bases his decision on the expected consequences, both material and psychological. First, each action results in

¹⁰This definition is identical to the one given in Köszegi and Rabin (2009).

different material utility in the final period. Second, unless the DM is certain about his health status in period 1, the remaining uncertainty is resolved in period 2. Thus there is a change in belief accompanied by gain-loss utility in the final period. The extent of this gain-loss utility is a function of the action chosen. Finally, if the DM decides to deviate from the planned choice there is an immediate change in belief accompanied by gain-loss utility in period 1. Suppose e.g. the DM has made the plan not to seek treatment. The expected utility of the two options “no treatment” and “treatment” are then given by:

$$\begin{aligned}
& \mathbb{E}_1 U^1(NT|NT) = \mathbb{E}_1 [u_2(NT)] \\
& = (1 - p_1)A + p_1B && \text{expected material utility in period 2} \\
& \quad - p_1(1 - p_1)\eta(\lambda - 1)r_{NT} && \text{expected gain-loss utility in period 2} \\
\\
& \mathbb{E}_1 U^1(T|NT) = \mathbb{E}_1 [u_2(T)] + \gamma_1 v(F_1|a = T, F_1|a = NT) \\
& = (1 - p_1)C + p_1D && \text{expected material utility in period 2} \\
& \quad - p_1(1 - p_1)\eta(\lambda - 1)r_T && \text{expected gain-loss utility in period 2} \\
& \quad + \gamma_1\eta[-\lambda(1 - p_1)\Delta_h + p_1\Delta_s] && \text{immediate gain-loss utility due to deviation}
\end{aligned}$$

Given the plan not to treat it is optimal not to treat if and only if $\mathbb{E}_1 U^1(NT|NT) \geq \mathbb{E}_1 U^1(T|NT)$. Let p_{NT}^* be the probability p_1 for which the DM is indifferent between the two options “treatment” and “no treatment” given he previously had the plan not to treat. Similarly, let p_T^* be the probability p_1 for which the DM is indifferent between the two options given he previously had the plan to seek treatment. As part of the proof of the following lemma we show that these probabilities are well-defined. Finally, define

$$V_T = \mathbb{E}_1 U^1(T|T) - \mathbb{E}_1 U^1(NT|NT) = -V_{NT}. \quad (\text{I.1})$$

We will see that the values $V_T(V_{NT})$ are an important part of the value of information. Denote by p^* the probability p_1 at which $V_T = V_{NT} = 0$. The three probabilities (p^*, p_T^*, p_{NT}^*) are all bounded away from zero and one, and have a clear order.

Lemma I.1. $0 < p_T^* < p^* < p_{NT}^* < 1$

Proof. See Appendix. □

The inequality shows that the probability of being sick below which it is optimal not to treat given the expectation not to treat, p_{NT}^* , is strictly larger than the probability of being sick above which it is optimal to treat given the expectation to treat, p_T^* . Hence, given a belief $p_1 > p_{NT}^*$ it is optimal to choose T irrespective of the planned

action. Equivalently, it is optimal to choose NT irrespective of the planned action if $p_1 < p_T^*$. This implies that the DM cannot rationally plan to take action T (NT) at information sets in which he entertains a belief p_1 that falls below p_T^* (exceeds p_{NT}^*). For information sets where $p_T^* < p_1 < p_{NT}^*$, he can plan to take the action he prefers at the planning stage as, once the plan is set up, he can expect himself to follow through on it.

From Lemma I.1 it follows that if the DM selects to test and his posterior probabilities are “extreme enough” his treatment choice must differ across information sets.

Corollary I.1. (1) *If $p^+ > p_{NT}^*$ and $p^- < p_T^*$, a plan $d_0 \in D_0^*$ prescribing to test (i), must also prescribe*

$$\begin{aligned} NT & \text{ if } s = s^- \\ T & \text{ if } s = s^+, \end{aligned}$$

i.e. an individual planning to test must plan to make his treatment choice dependent on the test result.

(2) *For all prior beliefs $p_0 \in (0, 1)$ there exist error rates $\epsilon^- > 0, \epsilon^+ > 0$ such that an individual planning to test must plan to make a signal-dependent treatment choice for all tests with smaller error rates.*

The corollary says that if the DM chooses to test and the posteriors are such that $p^+ > p_x^*$ and $p^- < p_y^*$ the DM will choose to seek treatment upon a positive test result and abstain from treatment upon a negative test result. Part (2) says that for any prior p_0 there exists a test that, if taken, influences treatment choice. Any test that is more accurate than this test will also influence treatment choice.

We can also make the following prediction concerning any plan that prescribes not to test.

Lemma I.2. *A plan $d_0 \in D_0^*$ prescribing not to test (n), must prescribe*

$$\begin{aligned} NT & \text{ if } p_0 < p^* \\ T & \text{ if } p_0 > p^* \end{aligned}$$

as continuation plan d_1 .

The reason is simple. Suppose it is part of a personal equilibrium to choose not to test. If this is true the DM expects to make the treatment choice based on the prior p_0 . Then by part (ii) of the definition of personal equilibrium he chooses the continuation plan, here the action choice, that gives him highest ex ante utility among all plans he will eventually follow. Thus, he will plan to seek treatment if and only

if $V_T = \mathbb{E}_1 U^1(T|T) - \mathbb{E}_1 U^1(NT|NT) > 0$ which is equivalent to $p_0 > p^*$. Lemma I.2 characterizes personal equilibria that involve test refusal when there is an option to test. It is easy to see that the result extends to the setting where there is no such choice. Lemma I.2 thus pins down what the DM would do if he could not test, assuming that in this case the DM would choose his preferred personal equilibrium. Hence we speak of action T (NT) being the DM's *default action* if $p_0 > (<)p^*$.

The considerations we have made so far focused on what treatment choices we can expect for a given belief p_1 . They are important when determining how a given test influences this choice. It is this potential influence that is a major part (in classic decision theory the only part) of an information's value. We will now turn to another consequence of choosing to test: the impact on the emotional response to learning the truth in the final period, or, more technically, the impact on contemporaneous gain-loss utility.¹¹

The Emotional Impact of Learning the Truth

Even if the test has no influence on the treatment choice, it still has an influence on utility in the final period. Unless the DM is certain about the state of health at the time he makes his treatment choice the remaining uncertainty is resolved in the final period. This results in a change in belief and thus triggers an emotional response. The information choice affects this final emotional response indirectly as it already resolves part of the uncertainty earlier. To give an example: if the test is perfect and turns out, say, positive the DM already knows with certainty that he is sick in period 1. Thus he cannot be disappointed anymore in the final period by learning about his bad health. More technically, as the information conveyed by the test changes the reference belief, the gain-loss utility from the evaluation of the final outcome against this reference lottery must also change. Given some treatment choice, this change in contemporaneous gain-loss utility (CGLU) due to choosing to test (i) instead of not

¹¹It is exactly such an impact of information choice that the literature of anticipatory utility neglects. Of course, anticipatory utility seeks to model "forward-looking" emotions, i.e. utility from beliefs about the future. Yet, if we think that our psychological well-being and, as an expression of this, utility is a function of our beliefs, why is it only a function of beliefs about some future event? It is equally reasonable to assume that living with a particular disease is psychologically at least as challenging as dreading this disease to break out in the future. If one, however, suspects uncertainty to be the driving force of dread, it needs to be underlined that this is not what is modeled by anticipatory utility.

testing (n) is

$$\mathbb{E}_0[v(F_2, F_1|i)] - \mathbb{E}_0[v(F_2, F_1|n)] = \begin{cases} q^+q^-\eta(\lambda-1)(p^+ - p^-)^2r_{NT} & \text{if NT chosen} \\ q^+q^-\eta(\lambda-1)(p^+ - p^-)^2r_T & \text{if T chosen.} \end{cases} \quad (\text{I.2})$$

The derivation is not difficult, but tedious. The results indicate that the DM gains from learning about his health status earlier apart from potential gains through making better decisions. This is because the information shifts the reference belief F_1 closer to the final belief F_2 , at least in expectation. This diminishes both gains and losses, but, as the DM cares more about the latter under loss aversion, the net effect is positive.

So far we have investigated two ways in which the DM gains (in expectation) from taking the test. We will now turn to the downside of testing: the expected emotional impact of information. This will turn out to be the major psychological cost of testing in this model.

The Emotional Reaction to the Reception of the Test Result

When patients test for a serious disease the reception of the test result usually triggers an emotional response. One advantage of the model of reference-dependent preferences proposed by KR is to offer a possibility to model this emotional response as the (dis-) utility derived from the change in belief triggered by the reception of the test result. The focus of this section is thus to derive the expected emotional impact of information, in formal terms $\mathbb{E}_0[v(F_1, F_0)]$.

The following observation will help the analysis. Under Bayesian updating the prior belief F_0 equals the expected posterior belief: $F_0 = q^-F_1^- + q^+F_1^+ = \mathbb{E}_0[F_1]$ where F_1^- is the posterior belief after a negative, F_1^+ the posterior belief after a positive test result. With this in mind, the binary signal structure allows us to make a useful simplification.¹²

$$\mathbb{E}_0[v(F_1, F_0)] = q^- \cdot \underbrace{q^+v(F_1^-, F_1^+)}_{v(F_1^-, F_0)} + q^+ \cdot \underbrace{q^-v(F_1^+, F_1^-)}_{v(F_1^+, F_0)}. \quad (\text{I.3})$$

The emotional impact of a negative result is thus a function of how the (factual) belief after these news compares against the (counterfactual) belief the DM would have had, had he received a positive test result. Similarly, the emotional impact of a positive test result depends on how the factual belief compares against the counterfactual. This observation is helpful as it allows us to confine attention to how the two posterior distributions F_1^- and F_1^+ compare against each other.

¹²It is important to note that this simplification is only valid due to the binary signal structure.

Note that the distributions F_1^-, F_1^+ are not only a function of the posterior beliefs p^-, p^+ but also of the treatment choice the DM plans to make in each contingency. The emotional response to information thus not only depends on the informational content of the signal but also on what one plans to do with this information. Let us distinguish two cases. First, call information *instrumental* if it affects the treatment choice, i.e., in this context, that the DM seeks treatment after a positive test result and abstains from treatment after a negative test result. Second, call information *noninstrumental* if it does not affect treatment choice. In that case the DM sticks to the action he would have chosen without any further information, the *default action*, no matter how the test turns out.

The Emotional Impact of Receiving Noninstrumental Information

The second case turns out to be the simpler of the two. If the information does not result in a change of action, both F_1^+ and F_1^- have the same support, the two material payoffs associated with the default action. The information received then only tells the DM how much probability weight to put on the high and the low outcome associated with the default action.

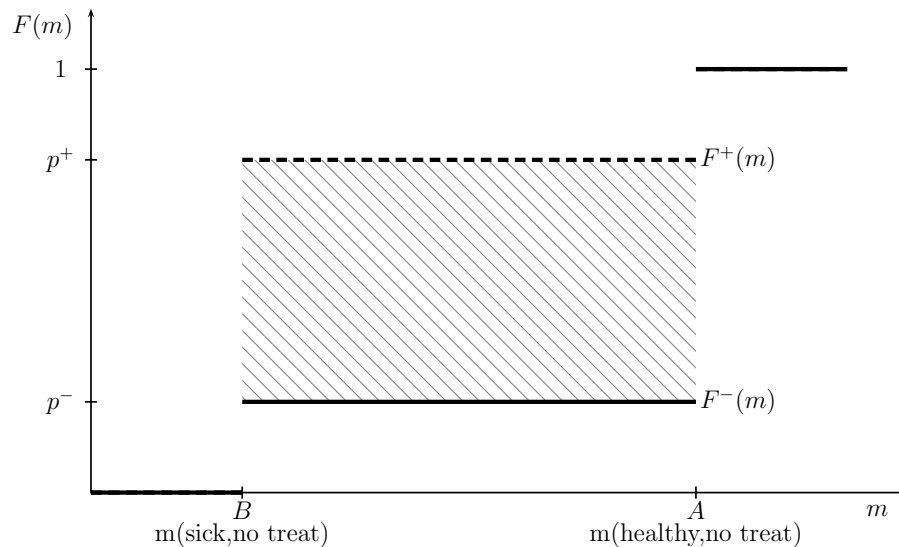


Figure I.1.: Difference between Posterior Beliefs when Information is Noninstrumental (default action NT)

Suppose e.g. the default action was NT (don't treat). The gain-loss utility derived

in period 1 (see Figure I.1) is

$$v(F_1, F_0) = \begin{cases} \eta [(p_0 - p^-)r_{NT}] = q^+ \eta (p^+ - p^-) r_{NT} & \text{if } s = s^- \\ \eta [-\lambda(p^+ - p_0)r_{NT}] = -q^- \eta \lambda (p^+ - p^-) r_{NT} & \text{if } s = s^+. \end{cases}$$

When information is noninstrumental, the emotions arising upon reception of the test result are a function of (a) the probability with which the alternative message was expected, (b) the distance between the posterior probabilities the two signals induce, and (c) the distance between the higher and the lower payoff associated with the default action. We can conclude that the emotional reaction is more severe, the more surprising the news (a), the higher the degree of certainty with which the news leave the receiver (b), and the more is at stake (c). Note that the emotional response to a negative test result is unambiguously positive and the emotional response to a positive test result is unambiguously negative.

Taking expectations over u_1 at $t = 0$ the expected utility in period 1 depending on one's default action is

$$\begin{aligned} \mathbb{E}_0 u_1 &= \gamma_1 [q^- v(F_1^-, F_0) + q^+ v(F_1^+, F_0)] \\ &= \begin{cases} -q^+ q^- \gamma_1 \eta (\lambda - 1) (p^+ - p^-) r_{NT} & \text{if } NT \text{ is default action} \\ -q^+ q^- \gamma_1 \eta (\lambda - 1) (p^+ - p^-) r_T & \text{if } T \text{ is default action.} \end{cases} \end{aligned}$$

The Emotional Impact of Receiving Instrumental Information

The case is different when the information is instrumental, i.e. when the DM seeks treatment after a negative and abstains from treatment after a positive test result. In this case, the distributions F_1^- and F_1^+ do not only differ in the probability they assign to being sick but also in the outcome associated with each state (see Figure I.2).

The gain-loss utility (PGLU) induced by the reception of a message is then

$$u_1 = \begin{cases} q^+ \gamma_1 \eta [-\lambda p^- \Delta_s + (p^+ - p^-) r_T + (1 - p^-) \Delta_h] & \text{if } s = s^- \\ q^- \gamma_1 \eta [p^- \Delta_s - \lambda (p^+ - p^-) r_T - \lambda (1 - p^-) \Delta_h] & \text{if } s = s^+. \end{cases}$$

The interpretation of this formula is easier if one remembers that this constitutes a percentile-wise comparison. A negative test result fares worse than a positive test result in the lower percentiles. This is because s^- induces no treatment and s^+ induces treatment, and the worst-case under a negative test (untreated and sick) fares worse than the worst-case under a positive test (treated and sick). On the other hand, a negative test result fares better in the middle and higher percentiles because (a) it puts more probability weight on the preferred state of the world (healthy) and (b) it results

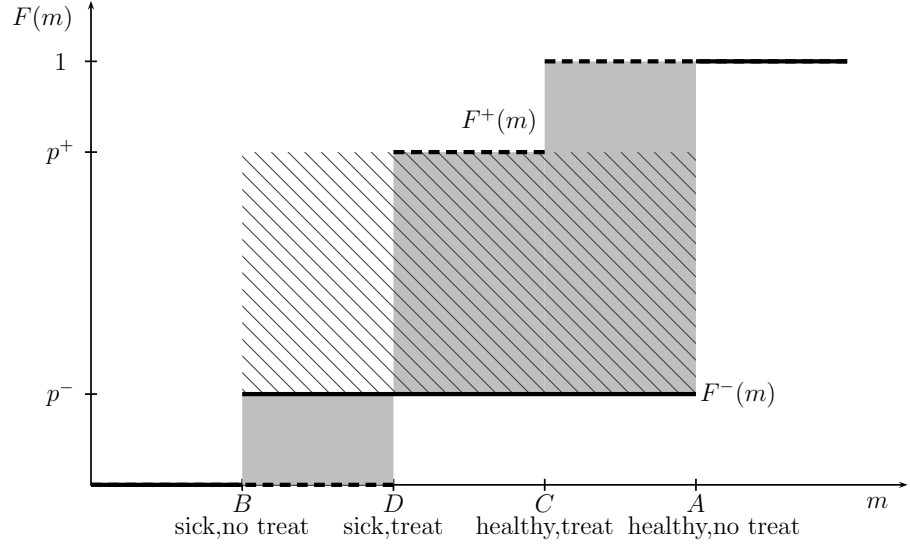


Figure I.2.: Difference between Posterior Beliefs when Information is Instrumental

in a higher payoff in the preferred state of the world. Taking expectations of u_1 at $t = 0$ the expected utility in period 1 is

$$\gamma_1 \mathbb{E}_0 v(F_1, F_0) = -q^- q^+ \gamma_1 \eta (\lambda - 1) [(1 - p^-) \Delta_h + p^- \Delta_s + (p^+ - p^-) r_T] \quad (\text{I.4})$$

$$= -q^- q^+ \gamma_1 \eta (\lambda - 1) [(1 - p^+) (\Delta_h + \Delta_s) - (1 - 2p^-) \Delta_s + (p^+ - p^-) r_{NT}] . \quad (\text{I.5})$$

Note the difference between the emotional impact of noninstrumental and instrumental information. While a negative result under noninstrumentality is unambiguously good it is accompanied by mixed feelings under instrumentality, unless the test is perfect. The same is true for the emotions generated by a positive result. While the effect is unambiguously bad under noninstrumentality it is accompanied by mixed feelings under instrumentality. Let us define the difference between the expected emotional response to instrumental and the expected emotional response to noninstrumental information, and call it *emotional differential*.

$$ED(NT) = -q^- q^+ \gamma_i \eta (\lambda - 1) [(1 - p^+) (\Delta_h + \Delta_s) - (1 - 2p^-) \Delta_s] \quad (\text{I.6})$$

$$ED(T) = -q^- q^+ \gamma_i \eta (\lambda - 1) [(1 - p^-) \Delta_h + p^- \Delta_s] \quad (\text{I.7})$$

Since the emotional impact of noninstrumental information depends on the default action, so does the emotional differential as the difference in emotional impact. The emotional differential will turn out to be one component of the value of information. The existence of this emotional differential suggests that instrumental information “feels different” than noninstrumental information controlling for the informational content.

Observation 1. *Beyond the content of information, the emotional response to information depends on whether the information is instrumental.*

We have now investigated the three channels through which the information decision affects utility. In the following section we connect the pieces to derive the overall value of information. This enables us to see when it is optimal to reject testing.

3. THE VALUE OF INFORMATION

The value of information (in utility terms) is computed by subtracting the expected utility of refusing to test from the expected utility of taking the test. To calculate the expected utility of each information choice we need to know what treatment choice follows. Lemma I.2 helps us to answer this question for someone opting out of testing. He will seek treatment if the prior p_0 is above p^* and he will not do so if the prior is below that threshold. The question concerning the treatment choice following the decision to test is essentially the question of whether the information is instrumental. As for some beliefs p_1 both treatment choices can be (credibly) made part of a plan, this question is far from trivial. We will proceed as follows. First, we derive the value of a test assuming that it is noninstrumental. Second, we derive the value of a test assuming it is instrumental. Finally, we apply the concept of preferred personal equilibrium to deduce which of the values is the appropriate one for a DM with a given prior p_0 facing the option of taking a test with given characteristics (ϵ^+, ϵ^-) . In order to illustrate the concept, the section concludes with an application to an analytically simple case: the value of a perfect test.

The Value of Noninstrumental Information

A test is noninstrumental if the default action is taken no matter how the test turns out. Formally, the value of noninstrumental information (VoNI), the difference in expected utility between the two alternatives of information choice $\{i, n\}$, is given by

$$\begin{aligned} VoNI &= \mathbb{E}_0 U^0(i) - \mathbb{E}_0 U^0(n) \\ &= \underbrace{\gamma_1 \mathbb{E}_0 [v(F_1, F_0)]}_{\text{PGLU}} + \underbrace{\mathbb{E}_0 [v(F_2, F_1|i)] - \mathbb{E}_0 [v(F_2, F_1|n)]}_{\text{Change in CGLU}} \\ &= \begin{cases} q^- q^+ \eta (\lambda - 1) [(p^+ - p^-)^2 - \gamma_1 (p^+ - p^-)] r_{NT} & \text{if } NT \text{ is default} \\ q^- q^+ \eta (\lambda - 1) [(p^+ - p^-)^2 - \gamma_1 (p^+ - p^-)] r_T & \text{if } T \text{ is default.} \end{cases} \end{aligned}$$

Note that the value of noninstrumental information does not depend on what information choice was planned if the test is noninstrumental, i.e. $\mathbb{E}_0 U^0(i|i) = \mathbb{E}_0 U^0(i|n)$ and

$\mathbb{E}_0 U^0(n|i) = \mathbb{E}_0 U^0(n|n)$. This is because a deviation from an anticipated information choice does not result in a change in the belief F_0 regarding future material payoffs.

One can make the following prediction concerning the desirability of noninstrumental information.

Proposition I.1. (*Kőszegi and Rabin (2009)*). *The value of noninstrumental information is positive if and only if*

$$\gamma_1 < p^+ - p^-.$$

As noninstrumental information is not affecting the treatment choice and thus material utility, its sole value lies in its potential to affect contemporaneous gain-loss utility, i.e. its potential to mitigate disappointment in the final period. This, however, comes at the cost of earlier disappointment when receiving the test result which is felt the more intense the larger the weight on prospective gain-loss utility γ_1 . The benefits of the information, meaning its potential to manage one's own reference point, increase in the distance between the potential posterior probabilities ($p^+ - p^-$). The larger this difference the more probable the reference point is shifted closer to the actual outcome. Lending the terminology of Eliaz and Schotter (2010) the test's value increases in the confidence with which it leaves individual.¹³ In line with empirical results provided by Eliaz and Schotter (2010) and Eliaz and Schotter (2007), the value of noninstrumental information (a) increases in the distance between the posteriors ($p^+ - p^-$) and (b) increases in the time distance between test results and outcome resolution if γ_1 is negatively correlated with this time distance as we assume.

The result stated in Proposition I.1 generalizes an example provided by Kőszegi and Rabin (2009) in which there is no action choice.¹⁴ If an information has no impact on subsequent decision-making it is sought if it is precise enough, but avoided if it is too imprecise.¹⁵ KR conclude that agents may seek to cluster information in order to receive one informative signal as opposed to a large number of less informative signals (avoidance of piecemeal information). However, as is pointed out in the next step, the more informative a signal, the more likely it is to affect behavior if optimal actions differ

¹³Instead of assuming a preference for confidence, i.e. making the difference between the posteriors an argument of the utility function, as is done in Eliaz and Schotter (2010) the RDP model endogenizes this "confidence effect" of information as a desire to manage one's reference point.

¹⁴KR assume a prior $p_0 = 1/2$ and posteriors $q, 1 - q, q > 1/2$. An interesting implication that is missed by this assumption is that, holding the quality of the signal fixed (i.e. its error rates), noninstrumental information loses value the more extreme the prior.

¹⁵This is an effect that is exactly opposite to the one predicted by models of anticipatory utility (see e.g. Barigozzi and Levaggi (2008)). With anticipatory utility, the value of noninstrumental information is strictly negative and decreasing in the signal's precision.

across states. But, instead of simply adding an instrumental value of information, the potential to affect behavior also influences an information's emotional impact. To see this, it is necessary to derive the value of an instrumental signal.

Value of Instrumental Information

Remember that a test is instrumental if a positive result leads to treatment and a negative result leads to no treatment. To determine the value of such a test one needs to compare the expected utility of each information choice. Similar to the value of noninstrumental information there is a term relating to the change in CGLU in the final period and a term capturing the emotional impact of information. In addition to these two terms, the value of instrumental information includes another two terms. First, the information choice affects material utility in the final period through its impact on treatment choice. Second, if the actual information choice differs from the planned information choice, there will be a change in belief at $t = 0$ resulting in prospective gain-loss utility (PGLU).

Take e.g. a DM who planned to take the test. The value of instrumental information (VoII) given by the difference in expected utility between the two possible information choices is

$$\begin{aligned}
VoII(i) &= \mathbb{E}_0 U^0(i|i) - \mathbb{E}_0 U^0(n|i) \\
&= \mathbb{E}_0 [u_2|i] - \mathbb{E}_0 [u_2|n] + \mathbb{E}_0 [u_1|i] - \mathbb{E}_0 [u_1|n] + u_0(i|i) - u_0(n|i) \\
&= \underbrace{\mathbb{E}_0 [m_2|i] - \mathbb{E}_0 [m_2|n]}_{\text{difference in expected material in t=2}} + \underbrace{\mathbb{E}_0 [v(F_2, F_1|i)] - \mathbb{E}_0 [v(F_2, F_1|n)]}_{\text{difference in expected CGLU in t=2}} \\
&\quad + \underbrace{\gamma_1 \mathbb{E}_0 [v(F_1, F_0|i)] - 0}_{\text{difference in expected PGLU in t=1}} + \underbrace{0 - \gamma_0 v(F_0(n), F_{-1}(i))}_{\text{difference in PGLU in t=0}}. \quad (I.8)
\end{aligned}$$

Similarly, given the plan of choosing not to test (n), the difference in expected utility between the two possible choices is

$$\begin{aligned}
\mathbb{E}_0 U^0(i|n) - \mathbb{E}_0 U^0(n|n) &= \mathbb{E}_0 [m_2|i] - \mathbb{E}_0 [m_2|n] \\
&\quad + \mathbb{E}_0 [v(F_2, F_1|i)] - \mathbb{E}_0 [v(F_2, F_1|n)] \quad (I.9) \\
&\quad + \gamma_1 \mathbb{E}_0 [v(F_1, F_0)] + \gamma_0 v(F_0(i), F_{-1}(n)).
\end{aligned}$$

Denote by W the value of instrumental information neglecting the term capturing the gain-loss utility in period 0 due to a deviation from the planned information choice.

$$W \equiv \mathbb{E}_0 U^0(i|i) - \mathbb{E}_0 U^0(n|n) \quad (I.10)$$

While $VoII$ captures the value of information at the moment of choice, W can be interpreted as the value of instrumental information at the planning stage. The latter

is important to determine the DM's preferences over plans. The former is important to determine whether a plan prescribing instrumental testing is self-enforcing, i.e. whether such a plan is a personal equilibrium.

The value of information W can be simplified to

$$W(NT^d) = q^+ V_T(p^+) + \text{VoNI}(NT^d) + ED(NT^d) \quad (\text{I.11})$$

$$W(T^d) = q^- V_{NT}(p^-) + \text{VoNI}(T^d) + ED(T^d) \quad (\text{I.12})$$

where $T^d(NT^d)$ denotes $T(NT)$ being the default action. The exact derivation is given in the Appendix. The equations show what determines the value of instrumental information. The second term gives the value of this test if it were noninstrumental. With reference to the discussion of the value of noninstrumental information we refer to it as the information's *confidence value*. This term is amended by the term $q^- V_{NT}$ (or $q^+ V_T$ respectively). It captures the test's impact on utility by changing the treatment decision away from the default action for one of the test results. We refer to it as the information's *instrumental value* as it is reminiscent of the classic economic idea of information deriving its value from influencing decision-making. Finally, the last term is the emotional differential capturing the difference in emotional impact stemming from the information being instrumental. The triple structure we find contrasts with the suggestion of Karlsson, Loewenstein, and Seppi (2009) that the value of instrumental information can be found by simply adding the instrumental value to the value of the information if it were noninstrumental. This procedure neglects the emotional differential between instrumental and noninstrumental information.

While above considerations focus on the part of the value of instrumental information that is independent of the planned information choice the value *at the moment of choice* includes a term capturing the (dis-)utility arising from deviating from the planned information choice. The value of information at the moment of choice is thus affected by which information choice $\{i, n\}$ was planned by the DM.

$$\text{VoII}(NT^d, n) = W(NT^d) + \gamma_0 \eta q^+ [p^+ \Delta_s - \lambda(1 - p^+) \Delta_h] \quad (\text{I.13})$$

$$\text{VoII}(NT^d, i) = W(NT^d) + \gamma_0 \eta q^+ [\lambda p^+ \Delta_s - (1 - p^+) \Delta_h] \quad (\text{I.14})$$

$$\text{VoII}(T^d, n) = W(T^d) + \gamma_0 \eta q^- [(1 - p^-) \Delta_h - \lambda p^- \Delta_s] \quad (\text{I.15})$$

$$\text{VoII}(T^d, i) = W(T^d) + \gamma_0 \eta q^- [\lambda(1 - p^-) \Delta_h - p^- \Delta_s] \quad (\text{I.16})$$

Equipped with these equations and the concept of preferred personal equilibrium we can now determine which of the equations actually applies to a given test.

The Value of a Test

Lemma I.1, I.2, and Corollary I.1 become very helpful when determining the value of a test with error rates (ϵ^-, ϵ^+) to a DM with prior p_0 . First, the two posterior probabilities (p^-, p^+) can be calculated via Bayes' rule. Second, the default action of the DM can be deduced using Lemma I.2.

There are two rather simple cases arising when the posteriors are extreme. First, if both posteriors are above p_{NT}^* (below p_T^*) the information must be noninstrumental as no plan prescribing the test to be instrumental would be followed through (Lemma I.1). Second, if $p^+ > p_{NT}^*$ and $p^- < p_T^*$ the information must be instrumental as no plan prescribing the test to be noninstrumental would be followed through (Corollary I.1). In all other cases, the difference between W and $VoNI$ determines whether the DM perceives the test as instrumental or not. We can thus distinguish four cases and their respective conditions for test refusal.

1. If $p^- > p_{NT}^*$, the only credible continuation plan, given that information is chosen, is to treat no matter how the test turns out. The value of the test is $VoNI(T^d)$ and ignorance is the PPE if and only if $VoNI(T) < 0$.
2. If $p^+ < p_T^*$, the only credible continuation plan, given that information is chosen, is not to treat no matter how the test turns out. The value of the test is $VoNI(NT^d)$ and ignorance is the PPE if and only if $VoNI(NT) < 0$.
3. If $p^+ > p_{NT}^*$ and $p^- < p_T^*$, the only credible continuation plan, given that information is chosen, is to treat if tested positive and not to treat if tested negative. The value of the test is given by $VoII$. Ignorance is the PPE if and only if $VoII(a^d, n) < 0$ and $W < 0$.
4. If $p^+ < p_x^*$ or $p^- > p_y^*$, there is more than one credible continuation plan $d_1 \in D_1^*$ given that information is chosen. The DM will prefer the test to be instrumental if $W > VoNI$. If the DM prefers the test to be instrumental, the value of the test is given by $VoII$ and ignorance is the PPE if and only if $VoII(n) < 0$ and $W < 0$. If the DM prefers the test to be noninstrumental, the value of the test is given by $VoNI$ and ignorance is the PPE if and only if $VoNI < 0$.

It is interesting to consider an analytically simple example that falls into the third category: the value of a perfect test ($\epsilon^- = \epsilon^+ = 0$). KR show that someone with $\gamma_1 < 1$ always prefers a perfect, but noninstrumental signal to remaining uninformed, and someone with $\gamma_1 = 1$ would be indifferent between receiving perfect information and remaining ignorant (compare Proposition I.1). In addition, perfect information

allows the individual to always make the right treatment choice. Perfect information must therefore have instrumental value. Taking these two considerations, intuition might lead us to expect that the value of information is strictly positive and that ignorance is never a part of the preferred personal equilibrium if the test delivers perfect information. It can be shown that there are cases in which this intuition is wrong.

The Value of Perfect Information

A test with error rates $\epsilon^- = \epsilon^+ = 0$ delivers perfect information about the state of health. The intent of this section is to discuss the value of such perfect information (VoPI). It can be shown that someone who is sufficiently confident to be healthy ($p_0 < p^*$) never rejects a perfect test while this might not necessarily be true for someone who entertains considerable doubts with regard to his health ($p_0 < p^*$).

Proposition I.2. *Preferences for perfect information. An individual with prior $p_0 < p^*$ never rejects perfect, instrumental information. On the other hand, for a person with $p_0 > p^*$, and $\gamma_1 > \frac{C-D}{A-D} = \frac{r_T}{r_T + \Delta_h}$ there exists a degree of loss aversion $\lambda^* < \infty$ above which this person prefers to ignore perfect, instrumental information.*

Proof. See Appendix. □

Now, why can it be optimal to reject perfect information for someone with T as default action while this cannot happen for someone with NT as default action? Both the *instrumental value* as well as the *value for confidence* of perfect information are positive if $\gamma_1 < 1$ regardless of the prior. The difference lies in the *emotional differential*. While it is strictly positive for someone with NT as default it is strictly negative for someone with T as default action. Remember that the emotional differential measures the different emotional impact of instrumental information compared to if it were noninstrumental. For someone with NT as default, good news feel worse but bad news feel better because of instrumentality. For someone with T as a default, bad news feel worse but good news feel better because of instrumentality. Due to loss aversion, the DM cares primarily about the emotional impact of bad news. Thus, making good news better while making bad news worse by an equal amount results in a more detrimental emotional impact in expectation and vice versa. If the patient with default T cares sufficiently about the emotional impact of information (γ_1 is large enough) and the degree to which he cares more about bad compared to good news is large enough (the degree of loss aversion λ is large), the rise in expected emotional disutility due to instrumentality outweighs the benefits of improved decision-making and higher confidence. It is then optimal to reject perfect information. This cannot happen to someone with default NT

because instrumentality implies improved decision-making as well as a more favorable emotional impact, i.e. a positive emotional differential.

It is still unclear whether people actually exhibit a degree of loss aversion that is high enough to make rejection of perfect information an optimal choice. There is some literature providing estimates of loss aversion. The results most commonly support estimates of λ of around 3, at least for monetary outcomes. A degree of loss aversion of 3 would be below λ^* assuming reasonable parameter values. Proposition I.2, however, illustrates the point that (a) the degree of loss aversion can have a positive or negative impact on the value of information depending on the prior belief p_0 and the extent to which the DM cares about the emotional impact of information γ_1 , and (b) the value of information does not necessarily increase when information becomes instrumental. While the second observation has also been made in the literature on anticipatory utility, the first observation points to a potential shortcoming of that literature. It has been shown by Eliaz and Spiegel (2006) that models of anticipatory utility are unable to accommodate information preferences that differ with the prior belief. Examples of such behavior are presented in Eliaz and Spiegel's article. Proposition I.2 indicates that the model of reference-dependent preferences may be able to accommodate such observations. The reason is simple and has been foreseen by Eliaz and Spiegel (2006): in the model we use here it is changes in beliefs that are important, thus utility depends both on posterior and prior beliefs.¹⁶

4. SCREENING AND THE TEST AS GATE-KEEPER

The section on the value of perfect information underlines that the instrumentality of information has two major effects on its value. First, there is a positive effect through improved decision-making. Second, there is an effect through a change in the emotional impact of information. While the example of perfect information illustrates the general working of the model it might not be appropriate for the cases of test refusal presented in the introduction as there exists no perfect test for most diseases. In addition, when looking at screening tests for HIV, colon, or breast cancer, patients should not expect to be treated without being tested positive. This is because (a) their prior probability of being sick is small, and (b) their action choice without a positive test result is restricted to "no treatment". The first reason underlines that the tests we are concerned about are screening tests, i.e. the test of patients that do not show symptoms of the disease. The second reason highlights that for diseases such as HIV or cancer positive tests

¹⁶Eliaz and Spiegel (2006) voice concerns with regard to adding prior beliefs to the utility function in addition to posterior beliefs. We comment on these in the conclusion of the chapter.

work as a gate-keeper to treatment. It is thus important to investigate under what circumstances the model can explain the refusal of an imperfect test by someone who is rather certain to be healthy and with “no treatment” as default action.

Asymptomatic screening is a type of medical test that is repeatedly a matter of discussion. On the one hand, it concerns testing patients that do not exhibit any a priori evidence of a disease as it amounts to testing a population that has a very low prior probability of being sick. On the other hand, there are diseases for which the benefits of identification of the disease are huge either because it facilitates containment of the disease (e.g. HIV) or it allows early treatment that comes with huge benefits compared to remaining untreated (e.g. breast cancer). It is often found that for diseases for which physicians recommend screening and patients generally acknowledge its benefits actual uptake rates fall below approval rates of patients.

This section seeks to replicate this setting in the following way. First, the patient’s prior belief p_0 is assumed to be close to zero. This and the fact that treatment is only possible after a positive test result make NT (no treatment) the patient’s default action. Second, we want to replicate a setting in which a physician recommends a screening test but the patient refuses. As a consequence the physician must face incentives different from the patient. We assume that the physician uses an expected utility model of decision-making while having a standard von Neumann-Morgenstern utility function given by m . This can be imagined as a situation in which the physician’s incentives are such that he seeks to maximize the patient’s expected health status $\mathbb{E}[m]$. He does not take into account the prospective or contemporaneous gain-loss utility arising from different decisions and health outcomes. This could be because he faces an incentive scheme (e.g. through the reimbursement/payment schemes of the medical system) that is purely focused on physical health outcomes. It could also be the result of the physician seeing himself primarily as a provider of physical health services leaving psychological services and considerations to specialized colleagues. The recommendation of the test is thus based only on its potential to improve expected physical health outcomes $\mathbb{E}[m]$ ignoring its psychological impact through gain-loss utility.¹⁷

Recall that the value of perfect information is strictly positive for someone with default NT . Thus ignorance can, if at all, only occur for tests with imperfect precision. However, to simplify calculations, we assume a false-negative rate $\epsilon^- = 0$ that implies $p^- = 0$. This assumption should not be a significant one in the screening setting as

¹⁷There are, of course, numerous other sources of incentive misalignment between patient and physician. In order to single out the effect of emotional responses to information we assume gain-loss utility to be the sole source of incentive divergence in this setting.

the prior p_0 is already assumed to be close to zero and $p^- < p_0$. Furthermore, if at all, it biases the results against ignorance as it increases the precision of the signal. Under these assumptions we can find the following condition for the rejection of a screening test.

Proposition I.3. *Rejection of Screening.* Assume an individual with RDP preferences and a physician using an expected utility-based decision model. Assume further that $p_0 < p^*$ such that without the test both agree on “no treatment” being the optimal action. Finally, assume that there is no false-negative error $\epsilon^- = 0$. The physician will recommend a test whenever the false-positive rate is such that $p^+ \geq \frac{\Delta_h}{\Delta_h + \Delta_s}$. Given a prior belief p_0 , there exists a range of tests, i.e. there exists a range of false-positive rates ϵ^+ , for which the physician recommends the test while the patient refuses to take the test if

$$\Delta_s < q^- r_{NT} \left(\gamma_1 - \frac{\Delta_h}{\Delta_h + \Delta_s} \right).$$

Proof. See Appendix. □

The condition for test refusal is interesting when comparing it to the reasons patients give for refusing a medical test. First, the lower the perceived risk the more likely is information refusal.¹⁸ The RHS is decreasing in p_0 , the subjective probability of being sick, through its impact on q^- making ignorance more likely the lower p_0 . Second, although recognizing the benefits of treatment, people reject information if treatment does not constitute a “cure”. As was explained in the model setup, the fraction $\Delta_h/(\Delta_h + \Delta_s)$ can be interpreted as the cost-to-benefit ratio of treatment. If benefits of treatment are large relative to costs, γ_1 will exceed this ratio making the RHS of the condition positive. In addition, $r_{NT} = r_T + \Delta_h + \Delta_s$, where r_T can be interpreted as a measure of effectiveness of treatment. If $r_T = 0$ treatment constitutes a perfect cure. Now, suppose r_T is large relative to $\Delta_h + \Delta_s$, i.e. there is a beneficial treatment which does not cure the disease (“care” instead of “cure”). This is the setting where it is most likely that the RHS exceeds the LHS, i.e. the condition for ignorance holds. Summing up, the setting in which the model predicts rejection of screening tests to occur despite recommendation by a physician are those in which a patient may test for a disease that (a) is severe, (b) for which there exists an efficient treatment but (c) this treatment is rather ineffective compared to a perfect cure. This seems to match the

¹⁸Note, however, that this is conditional on p_0 being close to zero, i.e., conditional on the individual to have low risk. As section 2.5 and I.2 indicate there is a stronger incentive to refuse information for high risks, i.e., those with $p_0 > p^*$. The model thus provides an explanation why, on the one side, one finds evidence for larger rejection rates among high-risk populations, yet, on the other side, not being at risk is a frequently mentioned reason for test decline.

situation with the medical screening tests presented in the introduction: the diagnosis of a predisposition for Alzheimer’s disease, for various types of cancer, or the diagnosis of HIV. There is a treatment available that is quite beneficial, but there is no perfect cure. Thus the model’s prediction is in line with the stated reason of “there is no cure” for information rejection.

5. COMPARATIVE STATICS

We are now interested in how certain parameters of the model affect the value of information and the incentives for test refusal. We look both on the impact of a variation in the parameters on noninstrumental and instrumental information. For simplicity, we look at the value of instrumental information W at the planning stage thus abstracting from the utility from deviation at $t = 0$. In addition, for instrumental information we confine ourselves to the value of perfect information. Concerning the question which parameters to vary, it helps to have a second look at the model setup where the interpretation of the different parameters is discussed.

First, we look at a variation in the health outcomes m . More precisely, we ask what happens if the benefits of treatment, the costs of treatment, and the severeness of the disease change. Second, we investigate what happens when the parameters γ_0 and γ_1 , that capture the impact of prospective gain-loss utility, are varied as these are assumed to be connected to the timing of testing. All proofs are in the Appendix.

Effectiveness of Treatment

An increase in the effectiveness of treatment (a) increases Δ_s , the net benefit of treatment, and (b) decreases r_T , that is the smaller the more effective the treatment.

Observation 2. *Treatment Effectiveness.* *If the test is noninstrumental, a change in treatment effectiveness has no impact on test uptake. If the test is instrumental, an increased treatment effectiveness increases the incentives for test refusal among those who would treat without further information but reduces incentives for test refusal among those who would (or could) not treat without being tested positive.*

It is important to highlight the difference of this result to the result obtained in a standard decision theoretic framework when assuming von Neumann-Morgenstern utility functions. In the latter framework the value of information is positive, thus information desirable, whenever the probability of sickness after a positive test p^+ is larger than $\Delta_h/(\Delta_h + \Delta_s)$. The question whether to obtain a test is, hence, a function of the accuracy of the test through p^+ on the one hand, and the efficiency of the treatment

through the cost-benefit ratio $\Delta_h/(\Delta_h + \Delta_s)$ on the other hand. A test can thus be desirable in a standard framework if treatment is highly efficient, though not necessarily highly effective. In contrast to this, a patient with reference-dependent preferences will be affected by treatment effectiveness, in addition to treatment efficiency, in his information decision. This becomes most obvious when considering the importance of treatment effectiveness for the rejection of screening tests as described in the previous section.

Costs of Treatment

The desire to be tested is also a function of the expected costs of treatment. A variation in these costs can be investigated while holding r_T and r_{NT} fixed. A rise in the costs of treatment increases Δ_h , but decreases Δ_s , the net benefit of treatment to a sick, when treatment costs are identical across states as is assumed.

Observation 3. *Treatment Costs.* *The value of noninstrumental information is independent of treatment costs. In contrast, if information is instrumental, an increase in treatment cost has a negative effect on test uptake for those who have “no treatment” as a default action. The effect on those with “treatment” as default action depends on the degree of loss aversion.*

One can, again, compare this result to the predictions of a standard decision-theoretic framework. As a rise in treatment costs decreases the efficiency of treatment, it leads to lower incentives for testing for someone with NT as default and higher incentives for someone with T as default. If patients have reference-dependent preferences the incentives in favor of testing for someone with NT as default decrease stronger in treatment costs and the incentives for testing for someone with T as default increase less in treatment costs compared to the predictions of standard decision theory.

Severity of the Disease

Finally, we are concerned whether the severity of the disease itself (r_{NT}) has an impact on the value of information. It is easy to show that this effect is rather simple.

Observation 4. *Severity of Disease.* *Holding the characteristics of available treatment, i.e. the benefit-to-cost ratio $(\Delta_s + \Delta_h)/\Delta_h$, and the relative effectiveness (r_T/r_{NT}) , constant, a variation in the severity of disease has no impact on uptake rates.*

Note that Observation 4 does not imply that the value of information is independent of the severity of the disease. It highlights that the sign of the value, and thus test uptake, is not a function of how serious a disease is, but of how likely it is, how the prospects (treatment potential) look like, and what kind of test is available.

Timing of Testing

KR suggest that the impact of gain-loss utility depends negatively on the distance between the point in time gain-loss utility is realized through a change in belief and the point in time the material utility is realized this belief is about. Denote by $\tau_1(\tau_0)$ the time distance between information reception in period 1 (information choice in period 0) and the resolution of health outcomes in period 2. KR thus suggest that $\partial\gamma_1/\partial\tau_1 < 0$ and $\partial\gamma_0/\partial\tau_0 < 0$. Taking up these assumptions we can address questions concerning the timing of testing.

First, consider a variation in the **speed of testing**. Fixing the time of information choice τ_0 , an increase in the speed of testing implies an increase in τ_1 . In addition, one can investigate the desirability of **earlier tests**. To vary the time of the test, without varying its speed, both τ_0 and τ_1 need to be varied simultaneously by an equal amount, say τ . While the effect of a variation in the speed of testing is straightforward, the effect of earlier tests is more ambiguous.

Observation 5. *Timing of Tests.*

(i) *The value of information, be it instrumental or not, is increasing in the speed of testing. Faster tests induce higher uptake rates.*

(ii) *If information is noninstrumental, earlier tests have higher value. Information rejection is less likely the earlier the test is conducted. If information is instrumental, earlier tests have a higher value if*

$$q^-(\lambda-1) [(1-p^-)\Delta_h + (p^+ - p^-)r_T + p^- \Delta_s] \left[-\frac{\partial\gamma_1}{\partial\tau_1} \right] > [\lambda p^+ \Delta_s - (1-p^+)\Delta_h] \left[-\frac{\partial\gamma_0}{\partial\tau_0} \right]$$

when “no treatment” is the default action, and

$$q^+(\lambda-1) [(1-p^-)\Delta_h + (p^+ - p^-)r_T + p^- \Delta_s] \left[-\frac{\partial\gamma_1}{\partial\tau_1} \right] > [\lambda(1-p^-)\Delta_h - p^- \Delta_s] \left[-\frac{\partial\gamma_0}{\partial\tau_0} \right]$$

when “treatment” is the default action.

The impact of an increase in the speed of testing is straightforward. As γ_1 is the weight on the emotional impact of information, that is, in expectation, always negative, the value of information always increases in the speed of testing. The lower the weight on prospective gain-loss utility in period 1, the higher the value of information. Thus, the faster the patient receives his test results, the more desirable is a test. The importance of the parameter γ_1 , and thus the speed of testing, is underlined by the fact that it plays a role in every condition for information rejection (see Propositions I.2 and I.3) we have found. Each of these conditions says that γ_1 needs to be sufficiently large to make ignorance the PPE. The model thus identifies the speed of testing as a crucial

factor in mitigating the emotional distress associated with information reception, and as a key variable in influencing test uptake.

The effect of the time of testing is similar to the effect of testing speed if information is noninstrumental, as for noninstrumental information there is no gain-loss utility in period 0. As $\frac{\partial \gamma_1}{\partial \tau_1} < 0$, the value of a noninstrumental test increases the earlier the test is conducted. As was already discussed in the section on the value of noninstrumental information this effect is in line with empirical findings by Eliaz and Schotter (2007).

Concerning the value of instrumental information, a change in the time of the test leads to an equally-sized change in τ_0 and τ_1 . There are two effects. The first one is familiar from the change in the speed of testing. It has a positive sign. The second term has a negative sign if the test's error rates are low enough, i.e. p^+ is large enough and p^- small enough. Assuming sufficient accuracy it is not straightforward to see which effect dominates. One can, however, make the following predictions. Suppose $\partial \gamma_0 / \partial \tau_0 = \partial \gamma_1 / \partial \tau_1$. If r_T is relatively large compared to $\Delta_1 + \Delta_2$ the first effect dominates. Thus, again, for diseases for which there is no perfect cure available, the second effect will dominate. Second, suppose γ_0 (γ_1) is a convex function of τ_0 (τ_1), i.e. the impact of gain-loss utility declines non-linearly in the time distance. Thus $\partial \gamma_0 / \partial \tau_0 > \partial \gamma_1 / \partial \tau_1$ which implies that the first effect is given a higher weight. Under these conditions, unavailability of a perfect cure and/or a higher responsiveness of the emotional impact in period 1 compared to period 0 to changes in timing, earlier tests are regarded with higher value by the patient. As earlier tests decrease γ_1 this results in higher uptake rates of tests.

These considerations are in line with observations concerning increased uptake rates for faster HIV tests.¹⁹ Patients seem to prefer faster tests even if the increase in speed comes at the cost of reduced accuracy. While this observation can also be explained by travel/time costs associated with collecting the test result when a second appointment is necessary, Observation 5 suggests an additional psychological motive for the observed preference for faster tests.

6. CONCLUSION

The chapter investigates how anticipated emotional responses to information affect information preferences. It seeks to highlight the connection between instrumentality and emotional impact of information and how they jointly determine the value of information. It is shown that instrumentality influences the emotional impact of information

¹⁹See e.g. Kassler, Dillon, Haley, Jones, and Goldman (1997) and Greenwald, Burstein, Pincus, and Branson (2006).

and this influence depends on an individual's prior belief. The impact of loss aversion on testing decisions may thus be quite different across people. Furthermore, it is possible to isolate different determinants of refusal rates. A large value of the parameter γ_1 , measuring the impact of prospective gain-loss utility, can be identified as a major determinant of test refusal. If the emotional effects of information play a large role one may look into ways how to dampen these effects. One possibility would be to address timing issues, in particular the speed of testing. It is also worth noting that uptake rates vary with treatment characteristics. Here, the availability of cure plays an important role. This is a result in contrast to the value of information in classic decision theory that does not depend on the size of potential benefits, i.e. on the gap between actual benefits and perfect cure. In this light it is worth remembering that it is "absence of a cure" that is stated as a rational for test refusal, and not that available treatments provide little benefit.

Eliasz and Spiegel (2006) raise concerns about incorporating prior beliefs in addition to posterior beliefs into the utility function in order to explain anomalies in informational choice. Incorporating priors in addition to posteriors into the consequence space is problematic from a revealed-preferences perspective as there is no choice problem that could reveal a preference over different priors. We agree. Yet, we find it important to note that this is not what the model of reference-dependent preferences we employ here does. In particular, the prior is not a part of the consequence space. The consequence space in the information choice remains to be described by posterior beliefs as in models of anticipatory utility. It is the preference order over this consequence space that depends on the prior.

With an ongoing debate on whether to transfer responsibility for medical decisions towards patients ("shared responsibility") it needs to be underlined that physicians' and patients' decisions can be quite different. If physicians use decision models based on expected utility theory thus not fully incorporating patients' preferences they will disagree with patients even if they are able to perfectly elicit the desirability of different health outcomes. The model presented here offers the patient's intent to balance material (or here: physical) and psychological concerns as a possible explanatory.

Finally, while we concentrate on the issue of evaluating medical information here, conclusions can be transferred to the evaluation of information in similar settings, i.e. to decision problems that exhibit a clear state preference on the side of the individual. It is worth investigating information preferences that result from reference-dependent preferences. With mounting evidence that a person's reference point is a function of his expectations, it is, at least partly, a result of prior information. Hence, when an individual's preferences are influenced by prior information acquisition and these

preferences determine future information acquisition, mutual interdependencies arise between information choice and information preferences. It would be worth investigating the results of these interactions in settings different from the one investigated here.

II. Product Lines, Product Design, and Limited Attention

We analyze how firms design their product lines when facing customers with limited attention. We assume that consumers simplify complex problems by neglecting some relevant aspects. Whether and to what extent a customer considers a particular attribute depends on the dispersion of that attribute in the set of alternatives. A firm may thus influence its customers' attention through the range of products it makes available. We show that a firm can increase its profit by introducing goods that have the sole function of manipulating consumer attention and discuss how a firm can profitably employ such manipulating goods.

1. INTRODUCTION

Most products have a whole set of attributes and features that are valued by the customer. For example, a car can be equipped with features like a sunroof, electric windows, or with a certain number of horsepower or safety features. Apart from these rather salient characteristics there are many others that are less noticeable. Examples would be the average durability of the gear box, the seating comfort, or the design of the cigarette lighter. Hence, such complex products like a car have a plethora of different attributes that all add to the overall quality of the product. Although information is freely available for lots of these attributes, consumers usually focus on a few key variables when making their purchase decision. Such systematic neglect of important information can arguably be classified as limited attention.

We propose a model of limited attention and analyze its implications for product design. As most consumer products have several characteristics, any purchase decision is a complex problem of trading off advantages in some dimensions against disadvantages in others. We posit that the way in which a decision-maker pays attention to different aspects of the problem reflects a need to simplify such complex decisions. The resulting attention allocation embodies both the decision-maker's valuation of the different aspects of a product, including its price, and the extent to which the available options (products and outside option) differ. We investigate how a firm optimally designs its product line when facing customers whose attention is determined in such a way.

We start with the optimal design of a single product that is offered by a monopolist. The optimal design reflects both customers' valuation and the production cost of a characteristic as in conventional models. In addition, the optimal design reflects the extent to which customers consider a particular attribute (the attention allocation) and how a particular design changes this attention allocation. The customers' tendency to focus on a few key characteristics creates an incentive for the firm to offer simple products. Such an incentive towards simplicity is interesting in light of a development of converging industries in which products include more and more features that substitute for previously distinct devices, such as cell phones including cameras¹, and an increasing debate over "feature fatigue".² Further, we find customers to profit from their inattention as it curbs the monopolist's ability to exploit its market power.

¹While such a development makes life more convenient for customers by replacing two previously distinct devices by one device, it adds difficulty by making that one device more complex.

²There is a discussion whether customers actually value simplicity or complexity. See e.g. Thompson, Hamilton, and Rust (2005) or Rust, Thompson, and Hamilton (2006). For an opposing view see e.g. Norman (2007). Our model of limited attention may add to this discussion a framework that allows to conceptualize a cost of complexity.

This is in strong contrast to the previous literature cited below that underlines firms' capabilities to exploit inattentive customers.

After considering the optimal design of a single product we investigate how a firm can profit from expanding its product line. We show that there is an incentive to offer more than one product to a set of customers with homogeneous preferences. This incentive to offer differentiated products stems from the possibility to influence the attention allocation by appropriately changing the choice set of the customers. One of the central propositions of the chapter is that firms will offer at least two kinds of products. There is a primary good that is intended to be sold to consumers. In addition, firms produce what we call bait goods. These are not intended to be sold but have the sole purpose of manipulating consumer attention in a way that is favorable to the firm.³ These bait goods tend to be high-end, state-of-the-art products that exploit technological boundaries. This finding may explain why shops usually put their most valuable, high-end merchandise on display. Optimal information provision or the minimization of search cost would suggest to put the product that is most commonly sold in the most prominent space. Yet, we see electronics shops featuring their largest plasma TVs and car dealers placing their most attractive sports cars in the most prominent shop spaces. Our model suggests that firms use such high-end products strategically to boost their customers' willingness-to-pay.⁴ Finally, we find that the optimal attention manipulation weighs the incentive to redirect attention to more profitable characteristics of the product against the incentive to maximize the attention that is paid to each characteristic of the product.

We join an emerging strand of the economic literature that incorporates limited attention into economic models. Gabaix (2011) develops a model of limited attention in which attention is the optimal solution to an attention allocation problem. Our approach draws heavily on his work, as the attention allocation we assume can be derived from an attention allocation problem with cognitive costs similar to the one proposed by Gabaix (2011). However, our modeling of the attention allocation differs from his approach as we employ some modified assumptions we deem appropriate to reflect the

³The assumption that a bait good is only used to manipulate consumer attention is useful to simplify the analysis. We later present an extension in which the bait good can also be profitably sold to some consumers. One example would be luxury goods that are sold to a small population of rich consumers and at the same time serve as a bait good to the larger population of less wealthy consumers.

⁴Similarly, Vikander (2010) reports a story about Audi advertising its \$118,000 R8 in the half-time of Super Bowl XLII. The advertisement spot cost Audi six million dollars. Since only a minority of Super Bowl viewers are able to afford such a car one may wonder why Audi did not choose other, far less costly marketing channels to reach potential customers.

costs of complexity. A more detailed comparison is deferred to later. In addition, the focus of our work are the implications of limited attention on market outcomes like product design and price. Kőszegi and Szeidl (2013) propose a model of focus-weighted utility that exhibits some parallels to our model. In an earlier version (Kőszegi and Szeidl (2011)), they show that a firm has an incentive to concentrate the advantages of its products in one dimension while spreading the disadvantages across as many dimensions as possible. While such an incentive is also present in our model (yet for different reasons, as we will explain in the next section), we derive more detailed results regarding the optimal product design. Zhou (2007) studies a monopolist's optimal product design if an advertising technology is available that highlights some of the product's characteristics. He investigates the potential and consequences of screening if customers are differently susceptible to such advertising. In our model, the way how customers allocate their attention is not determined by advertising, but by the attributes of the offered products. As this describes a more specific manipulation "technology", we are able to investigate the scope and limitations of such a manipulation. Also working on limited attention, Spiegel (2006a) and Spiegel (2006b) study the optimal industry behavior if consumers act according to the $S(1)$ sampling routine developed by Osborne and Rubinstein (1998). Beyond the attention heuristic employed, we deem the most important difference to be the welfare implications of limited attention. In Spiegel (2006a) and Spiegel (2006b) customers can be exploited by firms which obfuscate their products. Similarly, Rubinstein (1993) describes a firm's incentive to use complex pricing schemes to extract additional profits from boundedly rational customers. In contrast, we highlight that limited attention may primarily hurt the firm while benefiting customers, despite a firm's ability to manipulate attention. Eliaz and Spiegel (2011) propose a model that also features products with the sole function to attract attention. Consumers only consider the products of a subset of the firms in the market. Therefore, a firm uses attention grabbers if it wants consumers to consider its products. In contrast we investigate a firm's potential to attract or distract attention from product characteristics (including the price), thereby manipulating the desirability of a purchase. Bordalo, Gennaioli, and Shleifer (2012) and Bordalo, Gennaioli, and Shleifer (2013) develop a framework of limited attention to account for choice set effects. Their idea of limited attention is inspired by psychological findings concerning the perception of alternatives. These papers explain context effects through changes in the reference point that in turn influence the perception of different alternatives. Our framework models the impact of cognitive restrictions - not errors or biases in perception - on market outcomes. We will discuss the differences in more detail in the next section.

There is a large literature on choice set effects (Simonson (1989), Huber, Payne,

and Puto (1982)) and their impact on behavior in various settings (Herne (1997), McFadden (1999), Benartzi and Thaler (2002)). Several explanations for compromise effects⁵ have been proposed - ranging from extremeness aversion (Simonson and Tversky (1992)) to information inference from choice sets (Wernerfelt (1995), Kamenica (2008)). Kamenica (2008) shows that information inference creates an incentive to offer premium loss leaders. Though not explicitly relating his model to the compromise effect Vikander (2010) proposes a model of status effects and describes a firm's incentives to advertise premium products to an audience which is not able to afford the purchase. We show that in our model there is an incentive to offer premium products which are not necessarily intended for sale. Yet in our framework this incentive is based on a firm's ability to manipulate its customers' attention, and not on a firm's attempt to signal product value (Kamenica (2008)) or increase its products' prestige value (Vikander (2010)).

Johnson and Myatt (2006) describe how a firm may optimally design its product(s) to increase or decrease the dispersion of customer valuation. They find that a firm wants to concentrate its product's value in a single characteristic if a firm is confined to offer products with fixed expected value and there exists one characteristic for which customer tastes vary strongly. Furthermore, they investigate a monopolist's incentives to expand or contract its product line as the taste dispersion changes. In contrast, we discuss an incentive to expand the product line without any taste dispersion.

The remainder of this chapter proceeds as follows. In Section 2 we introduce and analyze the underlying attention process. Section 3 is dedicated to the derivation of the optimal product design if only one product can be supplied. This is extended in Section 4 where a firm can introduce an additional product that is designed to manipulate consumer attention. We continue by discussing some possible extensions in Section 5. Section 6 concludes.

2. THE ROLE OF ATTENTION

We assume a difference between *experienced utility* and *decision utility*. Experienced utility measures the satisfaction or welfare that customers derive from a choice. In contrast, decision utility depicts the way in which customers choose between alternatives. Decision utility thus not only depends on the welfare a customer derives from an alternative, but also on the way choices are made, and thus encompasses choice procedures, perceptions of alternatives (at the moment of choice), salience of attributes and alter-

⁵The compromise effect posits that expanding the choice set by a product which is more extreme in one attribute than any of the previously available options makes products which are mediocre in that attribute look more favorably.

natives, and the like. The distinction is supposed to depict the contrast between utility as a measure of welfare and utility as a tool to model choice (behavior).⁶ In our model the attention paid to attributes results in the difference between experienced utility and decision utility. The limited cognitive ability of humans to decide optimally on the plethora of information that is available to them forces them to focus on a subset of a problem's dimensions. We do not model any costs of obtaining information or costs of searching. Instead, we assume that all product information is readily available, but that consumers have problems of converting product information into an overall assessment of desirability. While we introduce the process of attention allocation and consumer choice and the underlying intuitions in the main text, we prove in the Appendix that the proposed attention allocation can itself be derived from an optimization problem.

The decision-maker faces a problem of deciding between a finite set A of alternatives $a \in A$. Each of these alternatives is described by a vector of attributes $(x_j^a)_{j \in J} \in \mathbb{R}^n$. J denotes the set of attributes according to which it is possible to distinguish between the alternatives $a \in A$, and $n = |J|$. Let experienced utility of an alternative $a \in A$ be denoted by

$$u(a) = \sum_{j \in J} v_j x_j^a.$$

The term $v_j \in \mathbb{R}$ measures the (marginal) value of attribute j to the decision maker (henceforth DM). The comparison of and choice between all alternatives in A is thus an n -dimensional problem.

The consumer's choice is based on decision utility however. Formally, the decision utility of an alternative is a function of the experienced utility of and the attention paid to each attribute:

$$\tilde{u}(a) = \sum_{j \in J} m_j v_j x_j^a. \quad (\text{II.1})$$

The term $m_j \in [0, 1]$ is the attention parameter associated with dimension j . If $m_j = 0$, dimension j is completely neglected. In this case, any differences between alternatives in dimension j will be irrelevant for the decision. We normalize the attention such that $m_j \leq 1, \forall j$.⁷ Since attention is normalized such that $m_j \leq 1$, an attention allocation of $m_j = 1$ means that the DM fully considers differences in dimension j when making her choice. Note here that we are back in the rational model if $m_j = m_{j'} > 0, \forall j, j'$.

⁶For a distinction of the concepts of experienced utility and decision utility, see Kahneman and Tversky (1984) and Kahneman (2000).

⁷Note that this normalization is without loss of generality since a decision utility with some attention vector \mathbf{m} yields the same choice behavior as a decision utility with an attention vector $\alpha \cdot \mathbf{m}$ for $\alpha > 0$.

To construct the attention allocation $(m_j)_{j \in J}$, let μ_j measure the maximal utility difference in dimension j between any two alternatives in the choice set A :

$$\mu_j = v_j \left(\max_{a \in A} x_j^a - \min_{a \in A} x_j^a \right).$$

We assume a complete ranking of dimensions according to μ_j . Let $r : J \rightarrow \{1, \dots, n\}$ be the function that assigns an attention rank to each dimension such that:

$$\mu_j > \mu_{j'} \Rightarrow r(j) < r(j')$$

where $r(j)$ denotes the attention rank of the attribute j .⁸ Having determined the attention order we now turn to the cardinal measure m_j of attention based on the attention hierarchy. Define:⁹

$$m_j = \max \left\{ 0, 1 - \frac{\kappa_{r(j)}}{\mu_j} \right\}.$$

The threshold $\kappa_{r(j)}$ is the minimum level of μ that a dimension with attention rank r needs to have in order to be taken into account. The threshold $\kappa_{r(j)}$ can be thought of as the cognitive cost of considering the r -th dimension of a problem. We make the following assumptions with regard to the thresholds κ_r :

Assumption II.1.

- (i) $\kappa_1 = 0$: *There is always one problem dimension that is fully considered.*
- (ii) $\kappa_r < \kappa_{r+1}$, $\forall r \in \{1, \dots, n-1\}$: *The attention threshold is increasing with each additional dimension that is considered.*

Let us give some motivation for our modeling of m_j . First, note that we do not model a problem of strategic attention allocation. While an optimization problem may, to some extent, underly the way in which attention is distributed, we assume it to be given at the moment of choice. Instead, we assume a particular rule how the attention parameters m_j are determined and give empirical and analytical reasons why this rule is sensible.

Note that we do not intend to model perceptual mistakes or biases. The DM is able to perfectly determine differences between alternatives in each dimension. The limitations in the cognitive process arise when the DM needs to integrate the information

⁸If there are dimensions j, j' that tie, $\mu_j = \mu_{j'}$, we assume the application of tie-breaking rules.

Assume e.g. $v_j \neq v_{j'} \forall j, j' \in J$ and $\mu_j = \mu_{j'} \cap v_j > v_{j'} \Rightarrow r(j) < r(j')$. The exact tie-breaking rule is of less importance. What is important, is that the attention ranking is a strict order, i.e. $\forall j, j' \in J, j \neq j' : r(j) \neq r(j')$.

⁹To let m_j always be well-defined let $m_j = 0$ if $\mu_j = 0$ and $\kappa_{r(j)} > 0$, and let $m_j = 1$ if $\mu_j = 0$ and $\kappa_{r(j)} = 0$.

about differences between alternatives in multiple dimensions with his own evaluation concerning the importance of each dimension in order to reach an evaluation of each alternative. This task necessarily includes making judgments as to how an advantage in one dimension trades off against a disadvantage or adds to an advantage in another dimension. Such judgments require the decision-maker to reach a conclusion with regard to the commensurability of different dimensions. For example, when deciding between different cars a decision-maker needs to make a judgment as to how an advantage of one car over another in terms of speed compensates for a disadvantage in terms of safety, or how it adds up to an advantage in terms of price, etc. As the number of relevant dimensions increases, so does the cognitive effort of evaluating all resulting trade offs. We argue that this need for commensurability judgments produces the complexity of a high-dimensional choice problem. We seek to model how a decision-maker deals with such complexity by simplification through neglect and prioritization of aspects of the decision problem.¹⁰¹¹

The attention allocation is to reflect a need to simplify a complex problem in order to be able to solve it. A sensible way to simplify is then to concentrate on those dimensions that are most important to the choice problem at hand. Therefore, the decision-maker will focus attention on those dimensions in which the largest utility differences occur: $\partial m_j / \partial \mu_j \geq 0$.¹² If the utility differences are sufficiently small, they are in a literal sense “negligible”, such that $\mu_j \leq \kappa_r(j) \Rightarrow m_j = 0$. However, since there are no commensurability judgments needed in a one-dimensional problem, and thus the DM faces a simple problem, there should be no difference between experienced and

¹⁰This might be the largest conceptual difference from other models of salience and limited attention.

In our model, the reason why the weightings m_j vary across dimensions is not that differences between alternatives are perceived as being larger or smaller than they actually are. In contrast, the weightings m_j express the extent to which differences between alternatives are appreciated in the decision process.

¹¹This is in contrast to previous literature on “choice overload” that highlights the number of alternatives as the basis of complexity and therefore suggests pruning the choice set to achieve simplification. If complexity is based in the necessity to make commensurability judgments between dimensions, as we argue here, pruning the choice set may not simplify the problem since as little as two alternatives can vary on a large number of dimensions.

¹²Empirical results on this issue differ. Results of Chetty, Looney, and Kroft (2009) indicate that consumers indeed often tend to pay higher attention to more dispersed attributes. They report that consumers are more sensitive to changes in the excise tax, which is included in the posted price, than to changes in sales tax, which is added at the register. As the posted price is larger than the additional tax that is added at the register, the two alternatives “buying”/“not buying” differ more in the posted price dimension than they differ in the additional tax-dimension. Thus, the more salient dimension seems to be the more dispersed one. For experimental results that challenge this assumption, see however Abeler, Falk, Goette, and Huffman (2011).

decision utility. This motivates part (i) of Assumption II.1. Incorporating more and more dimensions into the decision requires more and more commensurability judgments which becomes increasingly difficult. This is reflected in part (ii) of Assumption II.1. Furthermore, it motivates our assumption of a strict attention hierarchy. Allowing for dimensions to “share” an attention rank would result in the possibility that a decision-maker could perfectly solve a choice problem with arbitrarily many dimensions if it just so happens that $\mu_j = \mu_{j'}, \forall j, j' \in J$.

Consider an additional attribute when the complexity of the problem increases. This assumption implies that there is an attention hierarchy as it matters which attributes are considered “first”.¹³ Attributes that are ranked higher in the attention hierarchy have to meet stronger requirements for being taken into account. In addition, we assume that this increased complexity cost is reflected in a lower attention weight given that the attribute is in fact considered. Thus, differences in attributes that are ranked higher are appreciated to a greater extent.

One implication of assuming the ranking to be strict is that attention must vary between dimensions unless both are neglected ($\forall m_j > 0 : m_j \neq m_{j'}, \forall j' \neq j$). This is important because it implies that behavior is indeed distorted by limited attention. To see this, recall from the functional form of the decision utility (II.1) that if attention would be uniformly dampened (e.g. with $m_j = 0.5, \forall j \in J$), decision utility would just be a uniform transformation of experienced utility. In this case, a decision based on experienced utility is always the same as one based on decision utility. For limited attention to have a behavioral effect we need at least two dimensions j and j' which are allocated different levels of attention, i.e. $m_j \neq m_{j'}$. Our assumptions imply that any two dimensions that are considered receive a different level of attention.

Comparison to Other Models of Limited Attention

The proposed model of attention is conceptually very close to two models that have been proposed in the literature. We incorporate the characteristics of these models that we deem fit to depict the simplification process (and its cause) we have described above. In addition, we discuss where the models differ and why we deem these differences important. The first model on which we build is Gabaix (2011). We take over the idea of sparsity, meaning that some problem dimensions are neglected by the DM if they are not “important enough”. A major difference of our approach to the attention process of Gabaix is that the threshold $\kappa_{r(j)}$ associated with a dimension j depends on

¹³Note that the attention hierarchy is not meant to literally depict the timely sequence of how the DM takes attributes into account. Instead the attention hierarchy depicts how important an attribute is in the decision.

its *relative importance*, i.e. on how large its value and dispersion are compared to other dimensions' values and dispersions. While we retain the assumption that the vector of thresholds $(\kappa_1, \dots, \kappa_n)$ is exogenous, we impose a structure that we deem appropriate to reflect the notion of rising complexity costs as reflected in Assumption II.1. First, the decision-making process should not be distorted if the problem is not complex. Second, we argue that an increase in complexity should be reflected by an increased difficulty to consider more and more dimensions of a problem. This motivates our assumption of increasing thresholds $\kappa_{r(j)}$.

A second difference is the distinction between the attention, captured by m_j , and the valuation of a dimension, v_j . This distinction might seem superfluous at first glance as neither can be observed in isolation. Yet this simple distinction spares us the need to normalize our parameters to make salience independent of scaling. We find this feature desirable as we are concerned that some of the behavioral implications that Gabaix derives are based on this rescaling. In that sense certain behavior is predicted to occur not because people are inattentive but because this inattention is argued to be scale-independent.

Another model with which our model shares important characteristics is focus-dependent utility of Kőszegi and Szeidl (2013). They posit that the weighting of a dimension (m_j) is a strictly increasing function of the maximal utility difference associated with that dimension (μ_j). Our model shares this feature conditional on a dimension being considered, but not fully considered ($\mu_j > \kappa_{r(j)} > 0$). However, one major difference is that our model also allows a dimension to be fully neglected. We find this characteristic important as we see little sense in focusing attention if this does not entail a simplification of the problem to be solved. While the focusing model of Kőszegi and Szeidl (2013) (and also the salience model of Bordalo, Gennaioli, and Shleifer (2012) discussed below) may approximate zero attention weights for appropriately chosen parameters, there is a crucial difference between modeling neglect and approximating it. First, in all these models, including ours, the modeled biases result in a propensity to make judgment errors. Yet, in contrast to other models we point at a benefit of such a bias: the increased simplicity of the choice problem stemming from neglect. In Kőszegi and Szeidl (2013)'s model the focusing decision-maker solves a problem that is equally complex as the original problem yet is biased away from the original problem. Apart from these conceptual difference, an important difference between modeling and approximating neglect concerns their behavioral predictions. Suppose we find a decision-maker indifferent between two options that differ on a large number of dimensions. Models that only approximate neglect would predict a strict preference whenever one of the options is changed by only a very small amount in one

out of, say, a million dimensions. In contrast, our model predicts the existence of some dimensions in which small changes will remain unrecognized and thus will not result in the indifference to be broken. A second difference to the model of Kőszegi and Szeidl (2013) is the assumption of the attention hierarchy to be strict. This means that it is not possible to have two (or more) dimensions that receive the same, positive amount of attention m_j . Allowing for this to happen, Kőszegi and Szeidl (2013) predict unbiased choices in “balanced” choice problems.¹⁴ Their model thus predicts unbiased choices in arbitrarily complex problems, conditional on these problems being balanced. We deem such a model property to be at odds with our idea of how complexity problems arise and therefore insist on a strict attention hierarchy.

Bordalo, Gennaioli, and Shleifer (2012) and Bordalo, Gennaioli, and Shleifer (2013) present a model of salience-based choice. In contrast to our model, the modeling of salience is inspired by observations on perception (and not cognition) and assumes choice to be the result of multiple binary comparisons. More technically, the weighting does not only depend on the available choice set but also on the current binary comparison. Thus, an attribute’s salience in an alternative may be quite different depending on which two alternatives are currently examined. Bordalo, Gennaioli, and Shleifer (2013) applies the model to explain choice set effects such as the compromise effect. As others before¹⁵, he explains choice set effects by assuming that the choice set influences the decision-makers reference point, and thereby the evaluation of the individual alternatives. In contrast, our model does not involve the specification of a reference point. Choice set effects occur since the decision-maker focuses on a subset of the relevant problem dimensions, and the focusing procedure depends on which alternatives are available. Bordalo, Gennaioli, and Shleifer (2012) provide an interesting and important discussion how ordering principles governing the weighting of attributes, that are a main issue in our setting, may conflict with perceptual phenomena like diminishing sensitivity. In contrast, we assume the perception of differences to be given and concentrate on the obstacles arising from the need to integrate multiple perceived differences into an overall evaluation of desirability.

Closer to our understanding of attention is the model of rational inattention of Sims (2003). Like him, we do not seek to model costly information acquisition or information production (if one wants to understand perceptual biases in this way), but the

¹⁴In a balanced choice problem the two available options differ by the same amount of utility in each dimension in which one alternative has an advantage over the other, and by the same utility difference in each dimension in which this alternative has a disadvantage compared to the other. Utility differences must thus be equal across advantage dimensions, and across disadvantage dimensions, but may differ between advantage and disadvantage dimensions.

¹⁵See e.g. Drolet, Simonson, and Tversky (2000) and Kivetz, Netzer, and Srinivasan (2004).

problems associated with processing available information. Yet, while he focuses on the limits of data processing-capability, we seek to model the cognitive costs associated with solving a complex optimization problem.¹⁶

3. OPTIMAL PRODUCT DESIGN OF A MONOPOLIST

We want to investigate the problem of a monopolist designing a single product to be sold to customers with the described attention process. Suppose there is a set I of m qualities (save the price) that a product can have. Together with the price the purchase problem thus features up to $m + 1$ dimensions: $J = I \cup p$.¹⁷ For now, suppose that the level of each quality can take any non-negative real value: $q_i \in \mathbb{R}_+$, $\forall i = 1, \dots, n$ and $v_i \neq v_{i'}$, $\forall i \neq i'$. There is one attribute x_p that denotes the wealth of the decision maker. W.l.o.g. we normalize initial wealth to zero such that a value $x_p^a \leq 0$ means that alternative a is associated with a price of $P = -x_p^a$. Before we turn to the optimal design problem define a *null good* as an alternative with $q_i^0 = 0$, $\forall i \in I$.

A monopolist seeks to design a product that maximizes profit subject to the customer's willingness to buy it. To ensure that the customer is willing to purchase the good, the decision utility of the good (alternative a) must be weakly higher than the decision utility of abstaining from the purchase (alternative b). Note that alternative b is equivalent to a null good that is free of charge. Therefore, not buying is associated with a decision utility (and experienced utility) of zero. Part (i) of Assumption II.1 then implies that the monopolist cannot extract a positive profit by selling a null good at a positive price. Thus the product the monopolist designs must actually feature some qualities at positive levels in order to be sold.

Let the costs of producing quality level q_i be $c(q_i) = \frac{1}{2}c_i q_i^2$. The monopolist then maximizes his profit subject to the decision utility of alternative a being non-negative:

$$\begin{aligned} \max_{P, q_i} \quad & P - \frac{1}{2} \sum_{i=1}^n c_i q_i^2, \\ \text{subject to} \quad & \tilde{u}(a) = \sum_{i \in I} m_i v_i q_i - m_p v_p P \geq 0. \end{aligned}$$

¹⁶As an illustration of Sims' idea, think of a savings problem to which the optimal rule is to consume half of the income: $c_t = 0.5y_t$. Suppose income is a random variable and takes on the value 10.458376 at some t . In Sims' model the adaptation of c_t to the optimal value (5.229188) is costly as it requires the processing of the 8-digit input y_t . Note however that in Sims (2003), *finding* the optimal rule itself is not subject to cognition cost (though the anticipation of processing cost may alter the optimal solution itself). In contrast, we focus on the impact of limited attention on the *derivation* of the optimal solution.

¹⁷Note here that we differentiate between the n potential qualities of a product and the $n+1$ dimensions of the purchase problem.

Consider for a moment the case of unlimited attention $\tilde{u}(a) = u(a)$. It is straightforward to show that the optimal design then features all n qualities at levels $q_i = \frac{v_i}{v_p c_i}$ respectively. The monopolist reaps a total profit of $\Pi = \sum_{i \in I} \pi_i$ from the sale of the product, where $\pi_i = \frac{v_i^2}{2c_i(v_p)^2}$ denotes the profit from producing each quality i at its optimal level. We refer to π_i as the profitability of quality i in the following.

The following proposition states that under limited attention the optimal design is simpler in that it typically features only a subset $\mathcal{I} \subseteq I$ of all qualities. Denote by $\pi^{(t)}$ the t -th most profitable attribute (ties may be broken according to v_i).

Proposition II.1.

If a monopolist intends to supply a single good to the market, the optimal design features (i) only the most profitable quality ($i : \pi_i = \pi^{(1)}$) if and only if

$\nexists m \in \{2, \dots, n\} : \sum_{t=2}^m \left(\pi^{(t)} - \frac{1}{v_p} \kappa_{t+1} \right) \geq \frac{1}{v_p} \kappa_2$. In this case, the price extracts the whole surplus: $P = (v_i q_i) / v_p$.

(ii) Otherwise, it features the m most profitable qualities for which

$\sum_{t=1}^m \left(\pi^{(t)} - \frac{1}{v_p} \kappa_{t+1} \right) \geq 0$, while $\pi^{(m+1)} - \frac{1}{v_p} \kappa_{m+2} < 0$. In this case, the price does not extract the whole surplus: $P = \frac{1}{v_p} \sum_{i \in \mathcal{I}} m_i v_i x_i < \frac{1}{v_p} \sum_{i \in \mathcal{I}} v_i x_i$.

All qualities $i \in \mathcal{I}$ that the product features are produced at a level $q_i = v_i / (v_p c_i)$.

Intuitively, one can understand Proposition II.1 as showing how the limited attention of the customers translates into an augmented cost function of the monopolist. The limited attention reduces the willingness-to-pay for complex products, i.e. for products featuring more than one quality. This is equivalent to a monopolist facing additional costs when offering complex products. These costs include variable costs κ_{t+1}/v_p associated with introducing the t -th quality ($t > 1$), as well as fixed costs κ_2/v_p associated with offering a complex product in the first place.

A couple of things are noteworthy. Given the product features a particular quality, its optimal level is the same as the one under full attention. This is due to two effects resulting from the endogeneity of limited attention. First, limited attention reduces a firm's incentives to invest in quality as the created value is not fully taken into account by the customer. Second, the endogeneity of attention increases a firm's incentives to invest in quality as any additional unit of quality increases the attention paid to that quality and thus the decision value of any unit of quality already invested.¹⁸ In our model, these two effects cancel out each other perfectly, so the maximization problem of the monopolist yields the same level of product quality as in the case with unlimited attention for those qualities that are considered. However, the product

¹⁸Note that the firm's incentive to invest in quality in order to increase attention is conditional on the attribute being considered. If the attribute is not considered, small changes in quality might be insufficient in order to lift attention to a positive level.

generally features less qualities than under full attention and the price is lower if the product features more than a single quality. Note here that the exact canceling of the above-mentioned effects hinges on our modeling of limited attention, in particular our use of the range ($\max_{a \in A} x_i^a - \min_{a \in A} x_i^a$) as a measure of attribute dispersion. Yet, we claim that the second effect prevails as long as one retains the assumption of a higher attribute dispersion attracting attention. As this second effect compensates the first, limited attention does not necessarily imply the production of lower quality levels.¹⁹

We want to highlight the prediction of the design of simpler products, i.e., the impact of limited attention on the extensive margin of quality provision. It is a strong result in that it predicts certain qualities not to be produced at all despite there being a strictly positive marginal benefit and no marginal production cost at a quality level of zero. Previous work predicting such underproduction in the extensive margin has typically involved specific cost functions that resulted in the marginal production cost at the quality level of zero to be larger than zero.²⁰ In our framework, the tendency towards simpler products is entirely driven by the demand side, namely the complete neglect of certain qualities when customers derive their willingness-to-pay. This result is interesting in the light of the development of devices, e.g. smart phones, that include an ever-growing abundance of features. These devices certainly increase convenience as they combine the functions of several previously distinct devices. Above proposition questions whether this convenience will necessarily be rewarded with larger revenues since the increased convenience comes at a cost of increased complexity. Our model therefore suggests a way to conceptualize a value to simplicity in products. In addition, Proposition II.1 suggests a reduced form for modeling such incentives to the firm.

To motivate our next section recall that the monopolist cannot extract the whole value that the product yields to the customer as soon as the product features more than one quality. We will analyze a way in which the monopolist can increase the customer's attention to quality and thus the willingness-to-pay for the product that is sold to consumers without necessarily changing the product itself.

¹⁹In an earlier version to their 2013 article, Köszegi and Szeidl (2011) consider the implication of focusing for product design. Similar to above intuition, they find that qualities may be over- or underproduced depending on whether the attention function is concave or convex.

²⁰Köszegi and Szeidl (2011) e.g. find an incentive to concentrate the value of a product into a single dimension assuming a cost function that takes the sum of all quality levels $\sum_{i \in I} q_i$ as input. This can best be interpreted as a setting in which the firm may costlessly reframe the product such that different qualities are regarded as one cumulated quality or in which one quality may costlessly be split into several different qualities. Technically, it ensures that for all but one quality marginal production cost at a level of zero is strictly positive. The concentration result does not carry over to other cost functions, as e.g. the one employed here. This is because their focusing model does not feature neglect.

4. OPTIMAL PRODUCT DESIGN WITH INTRODUCTION OF A BAIT GOOD

Recall that the consumer attention is a function of both the valuation v_j and the dispersion within the choice set ($\max x_j - \min x_j$). A firm might not be able to change the valuation v_j of a quality, but it can manipulate the dispersion. If the firm can produce several goods, it might have an incentive to produce goods that have the sole purpose of increasing the dispersion of quality, thereby making these dimensions of the purchase problem more salient. Note that the quality levels are bounded below by the option of not buying ($\min q_i = 0$). This implies that in order to increase dispersion, products that are designed to manipulate attention must feature high quality levels. Such products will increase dispersion by increasing $\max q_i$. Yet, if these products are not intended for sale, they must be unattractive to consumers. This could, for example, be achieved by a very high price. This may however increase attention paid to the price dimension. Attracting attention to the price is not in the interest of the firm, as it makes a purchase less desirable to consumers. It may thus not always be in the interest of the firm to use such a manipulating device. In the following, we use the term *bait good* for those goods that have the sole purpose of manipulating consumer attention. Because these bait goods are designed to be unattractive, consumers still buy the main good that we henceforth call *primary good*.²¹ We investigate under what circumstances firm can profitably employ bait goods. We furthermore look at the characteristics of an optimally designed bait good.

In the analysis we assume that the firm does not incur any costs for designing and producing the bait good. We maintain this assumption to concentrate on the question whether it is possible to increase the willingness-to-pay of consumers for the primary good, and thus the profit made from its sale.

In addition, there are some arguments why the cost of the bait good might be negligible. First, if the bait good is never actually sold to the customer, it only needs a single item of the bait good that can be (unsuccessfully) offered to each of a large number of customers. If, as we are about to show, the bait good increases profits, the additional profits reaped from each customer may in sum be sufficient to cover the cost of bait good production. A second avenue to accommodate the cost of producing the bait good is to allow for customer heterogeneity in their intrinsic valuation of money. As we argue in the discussion section, a product designed for a richer customer segment is a prime candidate to serve as a bait good for a poorer customer segment. In this

²¹Obviously, the primary good still has to yield positive decision utility to the consumer in order to be bought.

way, the bait good is produced both for sale (to a rich customer segment) as well as for the manipulation of attention (of a poorer customer segment).

As outlined before, bait goods are designed to have premium quality. In reality, quality levels are typically bounded at some level. For example, there are technological and physical limitations on the horsepower that a car can have. Therefore, we now assume technical frontiers for each quality. Formally, all $q_i \in [0, \bar{q}_i]$ with $\bar{q}_i \in \mathbb{R}^+$.

Let us first look at the question under what circumstances a bait good actually increases profits of the firm. Note e.g. that if the monopolist is confined to producing a good with only a single quality, there is nothing to be gained from manipulating attention. This is because the primary good is designed such that the single quality employed ranks first in the attention hierarchy and thus receives full attention. This is different if the primary good features several qualities where, as we have shown before, the price gets full attention while the qualities do not get full attention. A bait good then derives its value through increasing the attention paid to the qualities that are not fully considered or are neglected altogether.

Proposition II.2. *Suppose a monopolist's profit-maximizing design of a single product features at least two qualities: $|\mathcal{I}| \geq 2$. If at least one of these qualities is not produced at the highest feasible level, the monopolist can strictly increase profits by using a bait good.*

Proof. See Appendix. □

The bait good thus increases the profit that can be made from the sale of the primary good while keeping the attention hierarchy among dimensions intact. This is done by designing the bait good such that it determines the range of available quality together with the outside option. The dual incentive to invest in quality that is present in the design of the single product, first to increase the value of the product given some attention by customers and second to increase attention of customers, is now split between primary and bait good. The bait good is designed to attract attention while the primary good is designed to provide value. To ensure this split, the bait good is made unattractive to the consumer with a sufficiently high price. This does not come at the danger of raising attention towards the price dimension if the optimal design of the primary good features at least two qualities. In that case, the price is already fully considered and people cannot pay more than full attention to the price.

Note here that we have only shown that a bait good can increase profits under fairly general conditions. We have not yet looked at the profit maximizing design of the bait good. One obvious venue to further increase profits is by increasing the qualities of the bait good as much as is technologically feasible. Note, however, that attracting more

attention to one quality may result in another quality receiving less attention since qualities “compete” over attention ranks. When contemplating the optimal design of the bait good the firm thus needs to make a decision as to whether it should exploit technological boundaries of a quality thus raising attention for an individual quality, but thereby potentially distracting attention from other qualities of the product.

We want to illustrate this trade-off inherent in the optimal design of the bait good in a setting with two qualities $\{1, 2\}$. We assume w.l.o.g. that quality 1 is more profitable: $\pi_1 > \pi_2$. We denote with $\nu_i = v_i \bar{q}_i$ the maximal utility that can be provided through quality i that is technologically feasible. We want to derive the optimal design of the bait good (q_1^b, q_2^b) which maximizes the profit from the sale of the primary good: $\Pi = \sum_i m_i^2 \pi_i$ with $m_i = \{0, 1 - \kappa_{r(i)}/q_i^b\}$. We will assume throughout that $q_1^b \geq q_1$, $q_2^b \geq q_2$ and $P^b > \sum_{i=1,2} v_i q_i^b / v_p \geq P$ such that the bait good is undesirable as it comes at an excessively high price. Note that this directly implies that $r(p) = 1$ for any design of the bait good. From the profit formula it is obvious that it is always optimal to maximize the level of the bait good in any quality as long as this does not change the attention rank of the quality. The question that interests us is which quality will feature a higher (lower) attention rank under the optimal attention manipulation. Let m_{iz} denote the maximal attention that can be attracted to quality i given that it ranks on position z in the attention hierarchy, assuming that attention is maximized for all higher ranking attributes given their rank.

Proposition II.3. *The optimal design of the bait good will attract more attention to quality 1 than to quality 2 if and only if $(m_{12}^2 - m_{13}^2)\pi_1 \geq (m_{22}^2 - m_{23}^2)\pi_2$.*

The proposition is straightforward from the profit formula. Yet, since the attention parameters are hardly observable we want to state sufficient conditions for quality 1 or quality 2 attracting most attention in a couple of corollaries.²²

Corollary II.1. *If $\pi_1 > \pi_2$ and $\nu_1 \geq \nu_2$, then the optimal bait good design will make quality 1 receive more attention than quality 2 and $(q_1^b, q_2^b) = (\bar{q}_1, \bar{q}_2)$.*

The intuition is rather simple. More attention can be attracted to the more profitable quality 1 at rank 2 than to the less profitable quality 2 at rank 2. In addition, placing quality 1 at rank 3 would restrict the extent to which the bait good can attract attention to quality 1. The bait good could not exploit the technological boundaries in quality 1. In contrast, placing quality 2 at the lower rank 3 does not restrict the bait good’s level of quality 2 to exploit technological boundaries and to set $q_2^b = \bar{q}_2$. We conclude that, if the more profitable quality is also the one with the larger technological boundaries,

²²See the Appendix for all proofs.

the bait good will be characterized by $(q_1^b, q_2^b) = (\bar{q}_1, \bar{q}_2)$. The bait good will be a state-of-the-art product featuring the highest possible quality in all dimensions.

Corollary II.2. *If $\pi_1 > \pi_2$, and $\nu_2 > \nu_1 > \kappa_3$, yet $\frac{\pi_1 - \pi_2}{\pi_2} \leq \frac{\nu_2 - \nu_1}{\nu_2} \frac{\kappa_1}{\kappa_2 - \kappa_1}$, then the optimal bait good design will make quality 2 receive more attention than quality 1 and, again, $(q_1^b, q_2^b) = (\bar{q}_1, \bar{q}_2)$.*

In this case, placing the more profitable quality 1 at the higher attention rank 2 restricts the extent to which attention can be attracted to quality 2. The bait good's level of quality 2 has to be restricted to a level below \bar{q}_2 for otherwise quality 1 could not maintain its higher rank. This again reduces the profit that can be made from quality 2. If differences in technological boundaries are strong enough and differences in profitability are rather small, it becomes optimal to attract more attention to the less profitable quality, and set $(q_1^b, q_2^b) = (\bar{q}_1, \bar{q}_2)$. Again, the optimal bait good is a state-of-the-art product. The crucial difference is that this premium product now attracts attention to the less profitable quality and distracts attention from the more profitable quality. The sufficient condition for this to happen relates the relative profit difference to the relative technological difference scaled by a constant factor.

In both cases considered so far, the bait good was a state-of-the-art product such that it featured the highest quality levels possible. This is not always the case.

Corollary II.3. *If $\pi_1 > \pi_2$, and $\nu_2 > \nu_1 > \kappa_3$, then for ν_1 sufficiently close to ν_2 the optimal bait good design will make quality 1 receive more attention than quality 2 and not exploit technological limits in all quality dimensions: $(q_1^b, q_2^b) = (\bar{q}_1, \nu_1/\nu_2)$.*

If the difference in technological limits is small enough, it is worthwhile to keep the more profitable quality at the higher attention rank. Yet, this implies that the extent to which attention can be attracted to the other quality must be restricted. The last corollary highlights how the trade-off underlying the optimal attention manipulation may produce different designs of the manipulating device, that is, the bait good.²³

5. DISCUSSION

In this work we seek to highlight how a firm may employ specially designed products to manipulate the attention allocation of its customers. There are a couple of issues and potential extensions of the model that are worth further discussion.

²³This trade-off is absent from models such as Kőszegi and Szeidl (2013) that do not feature the potential for attention distraction.

Heterogeneous Consumers

Until now we have assumed that consumers are homogeneous in both their preferences and their cognitive abilities. This assumption could be weakened in several ways. One interesting possibility is to allow customers to have different levels of wealth. For simplicity assume there are two groups of customers, very loosely labeled as *poor* and *rich*. Let the first group assign higher importance to the money dimension (and thus the price) of a product: $v_p^{poor} > v_p^{rich}$. This results in a firm designing different products to cater to both groups of customers. Recalling Proposition II.1, the rich group's ideal product may feature more qualities than the poor group's product. More interestingly, it features all the qualities the poor group's product has on a higher level. This makes the rich group's product an ideal candidate to be a bait good for the poor group. It increases the attention paid to all the qualities that the product for the poor features and thereby increases the willingness-to-pay of the poor. In this way, a firm may employ products designed for richer customer segments as a bait good for poorer customer segments. This can explain one of the observations we discussed in the introduction: the advertisement of expensive cars to an audience of which a majority is not able to afford it. It can also explain why shops tend to put their most expensive products, that can only be afforded by a small minority, on display.

If customers differ in the cognitive constraints they face, i.e. if they differ in their cognition costs κ_r , a firm may cater to these different groups with products differing in their degree of complexity. One reason for different cognitive constraints could be that one group has to act under stronger time pressure. If this is the case, the firm could, for example, offer products that differ in the number of qualities they feature. The hurried customer segment is then offered a product with only a few essential features, while customers with more time prefer more elaborate products.

Finally, customers may differ in their valuation for different qualities. As these valuations influence both a quality's profitability and the relative ease of attention attraction, a firm may have strong incentives to segment the customer population and design an appropriate product line for each segment separately. Still, analogous to the discussion of differing wealth levels, a firm may have an incentive to offer a product designed for one customer segment to a different segment despite that segment's unwillingness to purchase the product. Offering a sports car to a family father may increase his willingness-to-pay for horse power despite his general focus on car safety.

Negative “Qualities”

So far we have assumed that all qualities (save the price) are valued by the customer ($v_i \geq 0, \forall i$). There are certainly some characteristics that a product may feature which customers dislike, e.g. the level of exhaust fumes of cars, or the level of sugar and trans fats in food. While the customer dislikes these characteristics and their presence reduces consumers’ willingness-to-pay, a reduction and/or replacement of these attributes may be costly to the firm. Employing bait goods can thus have further beneficial effects since, in addition to attracting attention to positive characteristics, the firm could distract from negative characteristics. Also, a firm might have incentives to focus on reducing few negative characteristics a lot instead of reducing all negative characteristics a little bit. A label “sugar free” or “low carb” may effectively distract from other negative characteristics, for example high levels of trans fats.

6. CONCLUSION

We have introduced a new approach to model limited attention and applied it to the problem of optimal product design. The proposed attention heuristic fulfills several desirable properties that we think are realistic in real world markets.

Using this framework, we have shown that limited attention has far-reaching implications for product design and in general also for product lines. In general, a firm produces products that have fewer qualities than they would have if consumers were fully attentive. The customers’ tendency to neglect creates an incentive for the firm to offer simple products. Our model thus yields a way to formalize incentives to keep things simple.

We find that a monopolist would actually prefer consumers to be fully attentive. This is because the consumers’ willingness-to-pay is lower under limited attention since they do not fully appreciate all the qualities inherent in a product. Since the monopolist profits from an increase in attention, there is an incentive to introduce goods of premium quality that are not intended for sale, but that increase the attention of the consumers. These bait goods can be used to increase attention, but they can also distract consumer attention from attributes that are less profitable to the monopolist.

Although we have focused on product design, future research may apply the model of limited attention to a wide variety of settings. Given that we deem judgments of commensurability to be at the core of complexity problems, and our discipline being primarily concerned with the investigation of trade-offs across dissimilar dimensions, we are confident that our model may yield interesting insights into choice frictions in a lot of other contexts.

III. Limited Attention and the Demand for Health Insurance

This chapter presents an analysis of how customers with limited attention value and choose among health plans. We show how the model can accommodate four observations regarding plan choice. First, people tend to overweight the premium and thus underappreciate the value of health insurance. Second, insurance companies may have a strong incentive to reduce quality and to hide these shortcomings in the fine print while attracting customers with insufficiently lower premiums. Third, customers may choose dominated alternatives. Finally, the willingness-to-pay for insurance is subadditive creating an incentive for providers to unbundle comprehensive plans. We discuss how these effects may result in a fundamental dilemma for policy makers.

1. INTRODUCTION

Health insurance plans are complex products. They may differ in premiums, which diseases/treatments are covered, co-payment rates, deductibles, coverage of dependents, health incentives, maximum benefits, and many more aspects. Choosing whether to buy, and if so, which plan to buy, is thus seen as a difficult task. Accordingly, there is a strong need for advice on health plan choice and frequent efforts to provide such.¹ Advice on health plan choice and its many dimensions are also a frequent topic in the media.²

The topic of complexity of choice has attained increasing attention among economists. Specifically, the introduction of Medicare Part D prescription drug insurance in the United States offers an opportunity to study how consumers choose from a wide array of products differing on several dimensions. The question of whether consumers make optimal choices in that context has been the subject of extensive research.³ Both theoretical and empirical work has concentrated on the abundance of available options, thus highlighting the cognitive load associated with the number of options available.⁴ In consequence, discussion has focused on whether a restricted choice set would make consumers better off. This discussion however neglects a dimension of choice complexity. Choosing the right health plan is difficult not only because there are so many options available, but in particular because these plans differ in so many attributes.⁵ It is this aspect of choice complexity and its implications for choice among health plans that is the focus of this work.

This work highlights the difficulty associated with evaluating the desirability of a health plan given that each alternative may differ on several dimensions, in particular the numerous diseases for which treatment/medication may or may not be covered. It uses the model described in Chapter 2 to depict how consumers with limited attention choose among multi-dimensional options, by focusing on a limited number of the attributes that are relevant for choice. We assume that attention is directed towards those dimensions that incorporate the largest utility differences between the available

¹As an example, the U.S. federal government offers information on health plans and advice to understand the different features of health plans on www.HealthCare.gov.

²See e.g. <http://www.nytimes.com/2010/10/16/health/16patient.html>

³See Cubanski and Neuman (2007) and Neuman and Cubanski (2009) for reports on the program. See e.g. Heiss, McFadden, and Winter (2006) and Heiss, McFadden, and Winter (2010) for a discussion of optimal choice among Medicare Part D plans. Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2012) report evidence for a “comparison friction” in the choice of Medicare Part D plans.

⁴See e.g. Frank and Lamiraud (2008), Iyengar and Kamenica (2010), and Schram and Sonnemans (2011).

⁵One might argue further that the difficulty of having so many options would not arise if it was not for the possibility to vary health plans on so many dimensions.

alternatives. Dimensions in which the alternatives vary less receive less attention or no attention at all, i.e. they are neglected. The decision-maker thereby simplifies the problem before solving it.

We derive possible implications of such limited attention of customers. Specifically, we want to propose limited attention as a possible explanation for demand phenomena that have been observed in this context, yet are difficult to reconcile within existing models. First, we predict people to focus on the premium when deciding whether to purchase health insurance. The premium incorporates the largest utility difference across options for an average individual thereby attracting most attention in the insurance problem. This implies that people underappreciate or even neglect several of the benefits associated with having health insurance, and, consequently, undervalue health insurance. This is in line with empirical findings by Abaluck and Gruber (2011) as well as Heiss, Leive, McFadden, and Winter (2012). Second, the tendency of customers to focus on a limited number of diseases covered and to neglect the coverage of others allows firms to decrease the quality of their health plans unnoticed. We derive conditions under which customers choose health plans that exclude coverage for certain conditions not because coverage is deemed undesirable, but because the lack of coverage is ignored due to cognitive limitations. This topic has been mostly neglected in the health economics literature to date. Yet, media coverage and an abundance of Internet advice indicate that there is some popular interest in the topic of insurance providers hiding limited coverage in the fine print. Accordingly, policy makers have attempted to regulate information provision by insurance providers e.g. by restricting the use of fine print. In this chapter we argue that that it is not insufficient font size that results in customers making less-informed choices. We are thus skeptic about the effectiveness of policies banning fine print. Third, customers may make dominated choices if they happen to neglect exactly those plan attributes in which the domination occurs. This can explain the observation of dominated choices in the context of health insurance as documented by Sinaiko and Hirth (2011) that stands in sharp contrast to any preference-based model of choice. Finally, the undervaluation of comprehensive health plans creates an incentive for an insurance provider to unbundle and offer several more specific insurance plans individually. Such an unbundling effect has been reported in experiments by Johnson, Hershey, Meszaros, and Kunreuther (1993). We analyze which insurance plans are offered in market equilibrium and conclude the chapter by arguing that the effects we describe produce a serious dilemma for policymakers.

It has already been argued by Liebman and Zeckhauser (2008) that behavioral factors could play a major role in the markets for health care and health insurance. We seek to discuss one such factor, the complexity involved in the insurance purchase de-

cision. Similar to other authors, Liebman and Zeckhauser (2008) stress education and information provision as promising interventions to overcome behavioral biases. We fear that such interventions may be less promising with regard to the bias we discuss here. The complexity problems involved in the insurance decision are a result of the necessity to absorb and process a large amount information. Providing additional information might then turn out to be less helpful or even counterproductive. Close to our approach that assumes a necessity to focus on a subset of the available information is Kőszegi and Szeidl (2013). We are confident that some of our results can be replicated within the framework of Kőszegi and Szeidl (2013). This assures us that these predictions are not just an artifact of the model we employ here, but are robust to the choice of different models of limited attention. Distinct elements of our approach are the modeling of neglect and the assumption of a strict attention hierarchy.⁶ In particular, taking into account neglect enables us to accommodate evidence of dominated choices. Bordalo, Gennaioli, and Shleifer (2013) and Bordalo, Gennaioli, and Shleifer (2012) model salience as a function of perceptual biases. We share the assumption of the weighting of a problem's dimensions being driven by an ordering of dimensions with respect to utility differences although we differ on the motivation of this assumption. In addition to our approach, they discuss the conflicting influence of such an ordering and diminishing sensitivity in perception on the weighting process. Similar to Kőszegi and Szeidl (2013) their model neither features a strict salience order among the dimensions nor does it feature neglect. As discussed before, this makes their model incapable of explaining dominated choices. Customers' tendency to neglect and the resulting ability of firms to exploit such inattention have already been investigated in industrial organization settings (see most prominently Gabaix and Laibson (2006)). In our framework, inattention is endogenous. It has been shown that this endogeneity may lead to the opposite result of a firm being unable to exploit inattentive customers in a monopoly setting (see Chapter 2). Interestingly, we see the ability to exploit inattention reemerge when we look at more competitive markets in sections 3 and 6.

We will proceed as follows. Section 2 introduces the problem of whether to purchase health insurance and establishes the undervaluation of insurance if customers' attention is limited. Section 3 shows the great scope for profitably undercutting any incumbent health plan by reducing quality unnoticed. Section 4 discusses customers' propensity to make dominated choices. Section 5 summarizes experimental evidence suggesting

⁶We deem both elements to be important in the modeling of complexity problems as discussed in Chapter 2. Specifically, the assumption of a strict attention hierarchy ensures that more complex problems are more difficult to solve, thereby leading to a larger tendency for decision errors. In Kőszegi and Szeidl (2013)'s approach one can find arbitrarily complex problems which the decision-maker may solve without problem.

the profitability of unbundling comprehensive insurance plans and indicates how the model of limited attention can accommodate that evidence. Section 6 investigates which health plans are offered in market equilibrium. Section 7 concludes by shortly discussing some implications of these results for policy.

2. THE PROBLEM OF BUYING INSURANCE

We model the problem of buying health insurance as a problem of choosing between two multi-dimensional alternatives. Let the insurance contract be described by a premium P and a coverage rate $\alpha \in [0, 1]$, i.e. the share of treatment cost the insurance pays. Each alternative in the choice set is associated with a vector of consequences $((x_1, \pi_1), \dots, (x_m, \pi_m))$ where x describes the consequence, while π describes its probability. For example, buying health insurance incorporates the (certain) payment of a premium $(-P, 1)$. Second, it comprises consequences contingent on developing a disease. Let $I, |I| = n$ denote the finite set of diseases and assume that the occurrence of different diseases are disjoint events. For expositional purposes, suppose here that there are only two diseases $I = \{1, 2\}$. Buying insurance ensures the reception of treatment if need be. Thus it is associated with the health-related consequences $(-D_i + T_i, \pi_i)$, $i = 1, 2$ where D_i denotes the deterioration of health due to disease $i \in I$, $T_i \leq D_i$ the improvement of health due to medical treatment of disease i , and π_i the probability of developing disease i . Furthermore, dependent on α buying insurance is associated with a monetary consequence of a copayment in case of a disease: $(-(1 - \alpha)c_i, \pi_i)$ where c_i denote the monetary cost of acquiring treatment for disease i . In sum, the option *insurance* can be represented as⁷

$$((-P, 1), [(-D_i + T_i, \pi_i), (-(1 - \alpha)c_i, \pi_i)]_{i \in I}). \quad (\text{III.1})$$

In contrast, the alternative *remaining uninsured* is represented by a different vector of consequences. Suppose there are two types of diseases. First, there are diseases for which the decision-maker is able to afford treatment even when having no insurance. Denote this set by $F \subseteq I$ and suppose here $F = \{1\}$. Then in case of disease 1, having no insurance is associated with the health consequence $(-D_1 + T_1, \pi_1)$ since we assume the benefits of treatment to outweigh the cost. Second, contracting disease 1 is associated with the monetary consequence of having to pay the full cost of treatment $(-c_1, \pi_1)$. On the other hand, there may be diseases, for which the decision-maker is unable to afford treatment without insurance. Denote this set of diseases by

⁷For the moment, we assume that there are no diseases which are not covered by the insurance and no diseases for which the decision-maker is unable to afford treatment when having health insurance.

$\bar{F} = I \setminus F^8$ and suppose here that $\bar{F} = \{2\}$. For all $i \in \bar{F}$, remaining uninsured incurs the consequence of a deterioration of health $(-D_i, \pi_i)$. On the other hand, since the decision-maker is unable to afford treatment there is no monetary expenditure associated with receiving treatment, thus there is no monetary consequence.⁹ Thus the option of *remaining uninsured* is represented by a vector

$$([(-D_i + T_i, \pi_i), (-c_i, \pi_i)]_{i \in F}, [(-D_i, \pi_i)]_{i \in \bar{F}}). \quad (\text{III.2})$$

The decision-maker is both willing and able to purchase treatment for diseases $i \in F$. Yet, he is assumed to be willing but unable to do so for diseases $i \in \bar{F}$.

In order to solve the decision problem, the consequences of the available alternatives have to be ordered into categories, to which we henceforth refer as problem dimensions or aspects. One may think of this categorization process as a way of transforming the choice problem into a list of pros and cons of choosing one alternative over another. To arrive at such a representation, each consequence of an option is ordered into exactly one dimension. Two options each have a consequence in the same dimension if these two consequences are *comparable*.¹⁰ For example, the consequence of a co-payment in case of disease 1 associated with *insurance* is ordered into the same dimension as the full payment of treatment cost for disease 1 associated with *remaining uninsured* as they are both payments required to get treatment for disease 1. Similarly, the health consequence $(-D_2, +T_2, \pi_2)$ associated with *insuring* is categorized into the same dimension as health consequence $(-D_2, \pi_2)$ of a health deterioration associated with *remaining uninsured* as they both denote health consequences in case of disease 2. An option cannot have two consequences in one dimension. For example, the monetary consequences associated with disease 1 cannot be ordered into the same dimension with the monetary consequence associated with disease 2. Although both denote monetary consequences, they are associated with different events. Finally, there are consequences of one option that the other option lacks, such as the premium payment associated only with *insurance*, or the monetary expenses for treating disease 2, also associated only with *insurance*. This means that the option *remaining uninsured* lacks a comparable consequence that we represent by associating a consequence $(0, \pi)$ with that alternative

⁸Formally, if B is the decision-maker's budget, then $\bar{F} = \{i \in I : c_i > B\}$.

⁹Throughout, we abstract from monetary consequences associated with becoming sick that are unrelated to receiving treatment or insurance.

¹⁰The question of what constitutes a dimension of a choice problem is highly relevant for our later results since it is these dimensions among which attention is allocated. We assume a categorization based on comparability that we deem sensible for the choice problem under investigation. This is also consistent with our assumption discussed later that the difficulty associated with solving multi-dimensional problems is based on the need to make commensurability judgments with regard to different dimensions that are difficult to compare.

in the respective dimension. The categorization process leads to a choice problem with several dimensions. Denote by J the set of problem dimensions and note that the categorization described above leads to $|J| = 2n + 1$. The relevant aspects of choosing whether to insure comprise the payment of a premium (or the lack of it), as well as the monetary and health consequences associated with each disease $i \in I$.

The utility difference between two alternatives, here buying insurance versus not buying insurance, is assumed to be the sum of the utility differences in all dimensions of the choice problem. W.l.o.g. we assume the utility of an alternative in a dimension in which the respective alternative has no consequence to be zero.

Further, we assume the utility of a consequence to be linear. The difference U in utility between buying insurance and remaining uninsured is then

$$\begin{aligned}
 U &= v_p [-P - 0] + \sum_{i \in F} \pi_i v_h [(-D_i + T_i) - (-D_i + T_i)] + \sum_{i \in F} v_p \pi_i [(-(1 - \alpha)c_i) - (-c_i)] \\
 &\quad + \sum_{i \in \bar{F}} \pi_i v_h [(-D_i + T_i) - (-D_i)] + \sum_{i \in \bar{F}} \pi_i v_p [(-(1 - \alpha)c_i) - 0] \\
 &= -v_p P + \sum_{i \in F} \pi_i v_p \alpha c_i + \sum_{i \in \bar{F}} \pi_i v_h T_i - \sum_{i \in \bar{F}} v_p (1 - \alpha) c_i,
 \end{aligned} \tag{III.3}$$

where v_p denotes the marginal utility of money, v_h denotes the marginal utility of health. In our simplified setting with only two diseases this gives us

$$U = -v_p P + \pi_1 v_p \alpha c_1 + \pi_2 v_h T_2 - v_p (1 - \alpha) c_2.$$

The net utility of buying insurance comprises the disutility of the premium payment, the utility gain from receiving (partial) coverage of affordable treatment cost, the utility gain from receiving (otherwise unaffordable) treatment in case of disease 2, and the disutility from making a co-payment in case of disease 2.

The linearity assumptions on both the utility from consequences and the cumulative utility function U imply risk neutral preferences. Accordingly, if the premium is actuarially fair, i.e. $P = \sum_i \pi_i \alpha c_i$, the value of insurance U is given by

$$U = \sum_{i \in \bar{F}} \pi_i (v_h T_i - v_p c_i) > 0, \tag{III.4}$$

or $U = \pi_2 (v_h T_2 - v_p c_2)$ in our simplified setting. The utility of insurance is the net value of access to otherwise unaffordable treatment provided by the insurance. Following Nyman (2003) the value of insurance to a risk-neutral customer is created through the access motive: the possibility to acquire treatment for diseases $i \in \bar{F}$ for which the decision-maker is willing (thus the nonnegativity of the value) but unable to pay without insurance. We abstract from incentives due to risk preferences.¹¹

¹¹Our results do not hinge on the functional form of U since we an-

We suppose the decision-maker has difficulties thinking through each of the $(2n + 1)$ aspects in order to reach a decision. Specifically, incorporating every aspect of the problem into the decision requires judgments concerning the commensurability of different dimensions such as certain monetary consequences, uncertain monetary consequences, or uncertain health consequences. A full consideration of every aspect requires judgments as to how e.g. a disadvantage in the premium dimension compares to an advantage in terms of better health in case of a disease for which treatment is unaffordable. In addition, it requires a judgment as to how one advantage in e.g. a better health in case of one disease, say lung cancer, adds up to an advantage in better health in case of a different disease, say a flu. Appreciating each and every aspect of the problem thus requires a tremendous amount of commensurability judgments that all require cognitive effort. Thus, instead of fully evaluating the utility differences the decision-maker focuses on a subset of the relevant dimensions when making his choice. He thus bases his decision on the difference in decision utility \tilde{U} :

$$\tilde{U} := -m_p v_p P + \sum_{i \in F} m_{c(i)} \pi_i v_p \alpha c_i + \sum_{i \in \bar{F}} m_{h(i)} \pi_i v_h T_i - \sum_{i \in \bar{F}} m_{c(i)} \pi_i (1 - \alpha) c_i \quad (\text{III.5})$$

where $m_j \in [0, 1]$ is the attention a dimension $j \in J$ receives, $c(i)$ denotes the dimension comprising those consequences that refer to payments in case of disease i , and $h(i)$ denotes the dimension in which the health consequences of the options in case of disease i are compared.

Limited Attention

We want to shortly discuss the attention allocation that is reflected in the attention parameters m_j .¹² An implicit assumption in models of multi-attribute decision-making is commensurability, i.e. the possibility to measure and compare different concepts by a common standard. A decision-maker has no problem to determine how much a better treatment for disease i is worth compared to a worse (or no) treatment for disease j , or whether a better treatment for i is worth $\$x$ to her or not. The attention allocation we assume seeks to depict that the cognitive process of making different consequence dimensions commensurable, which is necessary to attain an overall assessment of desirability of one option over another, is difficult. It is this difficulty that makes complex

alyze a biased processing of the inputs of U . With $U = (1 - \sum_i \pi_i) [u(-P, -0) - u(-0, -0)] + \sum_{i \in F} \pi_i [u(-P - (1 - \alpha)c_i, -D_i + T_i) - u(-c_i, -D_i + T_i)] + \sum_{i \in \bar{F}} \pi_i [u(-P - (1 - \alpha)c_i, -D_i + T_i) - u(-0, -D_i)]$ one may model the familiar expected utility-representation of the problem. Assuming u to be concave over its first input (the sum of all monetary consequences in a state), one can model risk aversion. As it greatly simplifies exposition we opt for a risk neutral representation.

¹²A more extensive derivation and discussion can be found in Chapter 2.

decisions such as insurance purchase hard. The best way to simplify such decisions is to avoid the task of making dimensions commensurable by ignoring at least some dimensions. This may lead to worse decisions, yet it reduces the cognitive effort.¹³ It remains to ask which dimensions it is sensible to focus on. We assume that the decision-maker focuses on those dimensions in which the utility differences are largest. The decision-maker takes those dimensions into account in which the available alternatives differ the most, and, given these differences, the decision-maker cares most about. The attention allocation thus models neglect as the result of a process of simplification and prioritization. It endogenizes neglect by making assumptions about the characteristics of the dimensions that are ignored, instead of directly assuming ignorance with regard to specific dimensions. Finally, we assume this attention allocation to be “hard-wired”, thereby avoiding questions of strategic ignorance and infinite-regress problems.

We assume a particular framing of the choice problem. In this frame, the premium, the monetary consequences, and the health consequences of each individual disease form a distinct problem dimension. Why do we deviate from the familiar lottery representation, i.e. a problem representation based on states of nature? Our choice of a frame is necessary since we argue that the difficulty in complex problems is due to the necessity of making different dimensions commensurable. A frame different from the one we assume, in particular the lottery representation, already prerequisites this act of making different consequences commensurable. For one cannot compare the utility in a particular state (say disease i) between two alternatives without assigning a utility for this state to each of the two alternatives. Yet, this assignment already requires to integrate judgments concerning the relative desirability of different consequences such as a premium payment (or the lack of it), the health consequence, and a copayment (or the lack of one) into an overall assessment of the desirability of a particular alternative given that i occurs. Assigning the attention weights to different states instead of different consequences would thus assume that it is a comparison across states that is difficult and not the comparison of different consequences. That would contradict our very idea of what makes multi-dimensional problems complex.

We now want to formulate the attention weights m_j that reflect the above considerations. Let $\mu_j = \max_{g \in \Gamma} u(g, j) - \min_{g \in \Gamma} u(g, j)$ where $u(g, j)$ denotes the utility of that consequence of alternative g which is ordered into dimension j , and Γ denotes the set of all available alternatives g .¹⁴ For example, the utility of insuring associated with

¹³Note that considering not all of the alternatives, referred to as forming a “consideration set”, will not do the trick, since even as little as only two alternatives may differ on a large number of dimensions.

¹⁴Recall that we assume $u(g, j) = 0$ for those alternatives g with no consequence associated with dimension j .

the health consequence of disease 2 is $u(\text{insuring}, h(2)) = \pi_2 v_h(-D_2 + T_2)$ while the respective utility for remaining uninsured is given by $u(\text{uninsured}, h(2)) = \pi_2 v_h(-D_2)$. μ_j denotes the maximum utility difference in dimension j from any binary comparison of alternatives in the choice set. Since in the case we consider here there are only the two options of insuring and remaining uninsured, $\mu_{h(2)} = \pi_2 v_h T_2$.

We assume a strict hierarchy $r : J \rightarrow \{1, \dots, |J|\}$, among the problem dimensions $j \in J$ to which we henceforth refer as the *attention hierarchy*. This hierarchy obeys

$$\mu_j > \mu_{j'} \Rightarrow r(j) < r(j'), \quad (\text{III.6})$$

i.e., dimensions with larger utility differences attain a higher attention rank. In case that (III.6) does not produce a strict order we assume a particular tie-breaking rule.¹⁵ Given a dimension's rank in the hierarchy, the attention m_j it receives is given by

$$m_j = \max \left\{ 0, 1 - \frac{\kappa_{r(j)}}{\mu_j} \right\} \quad (\text{III.7})$$

where $\kappa_{r(j)}$ may be interpreted as the cognitive cost associated with considering the r th dimension of the problem. As we seek to model a decision-maker who has difficulties with solving complex problems we assume

$$(i) \kappa_1 = 0 \quad (\text{III.8})$$

$$(ii) \kappa_r < \kappa_{r+1}, \quad \forall r. \quad (\text{III.9})$$

This assumption reflects the rising difficulty of considering more and more dimensions of the problem. Eventually, if there are dimensions j , such that $\mu_j \leq \kappa_{r(j)}$, then $m_j = 0$. This means that any differences between the alternatives in these dimensions are completely neglected. The attention allocation thus reflects a need to simplify the complex choice problem in order to reach a decision. This simplification is achieved by ignoring some of the differences between the options. Due to limited attention the decision-maker may not (fully) appreciate differences between the two alternatives.

Undervaluation of Insurance

Returning to our problem of insurance purchase we consider the following assumption:

Assumption III.1.

$$v_p P > \{ [\pi_i v_p \alpha c_i]_{i \in F}, [\pi_i v_h T_i]_{i \in \bar{F}} \}. \quad (\text{III.10})$$

¹⁵If not stated otherwise we assume ties to be broken randomly.

The assumption states that (a) the premium exceeds the expected coverage of treatment cost of each *individual* disease, (b) for the diseases it provides access to treatment the value of the premium exceeds the expected value of this treatment for each *individual* disease. It turns out that this assumption is sufficient for the value of insurance to be underappreciated.

Proposition III.1. *Undervaluation of Insurance*

If Assumption III.1 is satisfied, then the decision-maker underappreciates the value of health insurance ($\tilde{U} < U$) that provides close to full coverage¹⁶ and may select not to insure despite it being individually optimal.

Proof. See Appendix. □

We argue that this implies that a majority of people underappreciates the value of health insurance. Consider the setting for which Assumption III.1 is satisfied. Part (a) of the assumption is always satisfied if the premium is greater or equal the actuarially fair premium, i.e. if the insurer breaks even. Part (b) is satisfied if insurance covers sufficiently many diseases that are unlikely *individually*. We regard this assumption to be satisfied in the case of health insurance for the average customer, i.e. an individual with no severe pre-existing condition. This is based on the observation that the distribution of medical expenditures is highly skewed. The most common diseases have rather minor health consequences and available treatments tend to be cheap. These should constitute the diseases in F . On the other hand, the largest chunk of medical expenditures is created by rare diseases with highly expensive treatments. These should be the ones we would expect to form the set \bar{F} . Hence, we regard it as a valid approximation to assume that the probability to contract any particular disease $i \in \bar{F}$ is small. If (and only if) Assumption III.1 is satisfied, the premium dimension receives full attention, $m_p = 1$, while all further dimensions $j \neq p$ are not fully considered, $0 < m_j < 1$, or even neglected, $m_j = 0$. Note that this implies that the cost of insurance (the premium) is fully considered while its benefits are not fully appreciated.¹⁷ Consistent with this, Abaluck and Gruber (2011) find that elders place too much weight on the premium relative to expected out-of-pocket costs when choosing a Medicare Plan D prescription drug plan.¹⁸

¹⁶Precisely, if the premium P is nondecreasing in the level of coverage α , and the premium for full coverage $\alpha = 1$ is affordable, then there exists a level of coverage $0 < \underline{\alpha} < 1$ such that the decision-maker underappreciates the value of insurance for any health plan (α, P) with $\alpha > \underline{\alpha}$.

¹⁷The necessary and sufficient condition for full insurance to be underappreciated is $m_p > \bar{m}$, where \bar{m} is the weighted average of the attention parameters associated with the dimensions $j \neq p$.

¹⁸In addition, they find elders not to value variance-reducing aspects of health plans. This last finding is particularly striking as variance reduction is the classic argument for insurance purchase. We

The value of insurance U offering (close to) full coverage is not fully appreciated.¹⁹ First, people tend to underappreciate all the out-of-pocket cost of attaining treatment that health insurance takes over. Intuitively, as the number of potential diseases is large people are unable to take into account all the expected cost they have to cover privately if they remain uninsured. Second, people tend to underappreciate the access value provided by health insurance. Again, as the number of potential diseases is large, the decision-maker is unable to consider each disease for which he will not be able to afford treatment if remaining uninsured.

Next to insufficient income to afford premiums or alternative ways to access medical care (e.g. charity or Medicaid in the U.S.) this underappreciation of the value of health insurance can explain the prevalence of a significant number of voluntarily uninsured where health insurance is not mandatory. Consistent with the model's predictions, Heiss, Leive, McFadden, and Winter (2012) report an undersubscription to the generous Gold plans compared to the Silver plans with less benefits and lower premium. Although being primarily concerned with the sources of advantageous selection, Fang, Keane, and Silverman (2006) find cognitive ability to be positively correlated with insurance purchase. We predict this under the natural assumption that the cognitive cost parameters κ_r are negatively correlated with cognitive ability.

Our result of an undervaluation of insurance may be contrasted with evidence suggesting a preference for excessively low or no deductibles.²⁰ More generally, economists have consistently argued that people tend to overinsure against health risks.²¹ First, we want to emphasize that the characterization of a preference for low or no deductible or a strong preference for full insurance as overinsurance is based on the common approach to restrict the value of insurance to the balancing of consumption across states, i.e. its risk-reducing function. If one assumes the value of health insurance to be primarily driven by access motives, as we do here, a preference for a low deductible or even for full insurance cannot be understood as overinsurance. A deductible as high as \$1,000 may already restrict access to medical care if a household falls on hard times. Choosing a

conclude that our approach to disregard incentives based on risk aversion can be viewed as a reasonable approximation.

¹⁹We conjecture that underappreciation holds much more generally. For example, if α is close to zero, the set of diseases for which insurance provides access to treatment (call it A) is empty, i.e. the health insurance provides no access value, and $U - \tilde{U} < 0$. We cannot completely rule out the possibility that $U - \tilde{U} \geq 0$ for intermediate α for all possible $(v_p \alpha c_i)_{i \in F}$, $(v_h T_i)_{i \in A}$, $(-v_p(1-\alpha)c_i)_{i \in A}$ with $A = \{i \in I : B < c_i \leq (B - P)/(1 - \alpha)\}$. Yet, we conjecture that the underappreciation of coverage and access value usually dominates the underappreciation of copayment.

²⁰See e.g. Sydnor (2010).

²¹See e.g. Feldstein (1973) and Feldman and Dowd (1991).

low deductible or no deductible at all ensures this access.²² We want to underline that we do not seek to negate the role of risk preferences for the demand of health insurance. Yet, given decades of finding evidence that economists interpret as “overinsurance” we think that we should at least consider the possibility that it is not the customers who consistently buy too much insurance, but it is us economists neglecting an important part of the value of insurance when deriving the optimal amount of insurance.

The underappreciation of the value of insurance is a result of the complexity of the insurance-purchase problem. Such underappreciation does not only make the option of remaining uninsured more attractive as it actually is, it also makes insurance plans with lower coverage more attractive as they are. This gives rise to a strong potential for undercutting.

3. PROFITABLE UNDERCUTTING IN AN INSURANCE MARKET

In this section we want to show that an insurance provider entering the market may profitably attract customers from an incumbent insurance plan by undercutting the premium and lowering coverage. We do not yet consider a full characterization of firm behavior in the insurance market. This will be addressed in a later section. Here we seek to establish that customers with limited attention are attracted towards low-premium, low-quality plans to a suboptimally strong degree and that firms may exploit this attraction. We argue that it is this exploitation of limited attention that underlies the frequently-voiced suspicion that firms “hide” shortcomings of the products in the fine-print of the contracts.

Let us consider more generally the undercutting strategy described above. First, assume that there is an incumbent insurance plan, e.g. a public insurance program. Let it be characterized by some premium P and some degree of coverage of health costs $\alpha \in (0, 1]$. Define $A(\alpha) \subseteq \bar{F}$ as the set of diseases for which an insurance plan with coverage rate α provides access to treatment.²³ To save on notation, let $A = A(\alpha)$ denote the

²²A different explanation for a preference for low deductibles would need a modification of the model we apply here. Suppose the explicit mention of a deductible increases the salience of exactly those instances in which the insurance does not pay. Further suppose, that the attention rank of a consequence does not only depend on the utility difference across alternatives but also on the salience of the consequence. In this case, the decision-maker will focus on the events in which a high-deductible insurance does not pay off while neglecting the ones in which it does. A high-deductible insurance may then be regarded as receiving (close to) no insurance yet with the obligation to pay a premium.

²³Formally, $A(\alpha) = \{i \in I : B < c_i \leq (B - P)/(1 - \alpha)\}$. We assume throughout that A is nondecreasing in α in the sense that $i \in A(\alpha') \Rightarrow i \in A(\alpha)$, $\forall \alpha' \leq \alpha$. This will hold as long as $(B - P)/(1 - \alpha)$ is nondecreasing in α .

set associated with the incumbent plan. Furthermore, assume that the premium is at least actuarially fair, i.e. $P \geq \sum_{i \in F \cup A} \pi_i \alpha c_i$. Now, consider a second insurance plan with $0 < \alpha' < \alpha$ and denote by $A' = A(\alpha') \subseteq A$ the set of diseases for which this second health plan provides access to treatment. Assume that this second plan is priced such that the premium difference reflects the difference in expected cost of coverage, i.e. $P' = P - \sum_{i \in F \cup A'} (\alpha - \alpha') \pi_i c_i - \sum_{i \in A \setminus A'} \pi_i \alpha c_i$. The lower premium P' incorporates the saving of expected coverage cost for all treatments that will be demanded by customers under both plans $i \in F \cup A'$. In addition, there may be additional cost savings as a reduction in coverage may make some treatments unaffordable as copayments $(1 - \alpha)c_i$ increase. For these diseases $i \in A \setminus A'$, the insurance does not have to cover any treatment costs. We now want to investigate the difference in decision utility between health plan 1 and health plan 2. If this difference is negative the second, low-coverage health plan is preferred to the first, high-coverage health plan. The difference in decision utility is given by

$$\begin{aligned}
\tilde{U}_1 - \tilde{U}_2 &= m_p v_p (P' - P) + \sum_{i \in A \setminus A'} \pi_i [m_{h(i)} v_h T_i - m_{c(i)} v_p (1 - \alpha) c_i] \\
&\quad + \sum_{i \in F \cup A'} \pi_i m_{c(i)} v_p (\alpha - \alpha') c_i \\
&= m_p v_p \left[\sum_{i \in F \cup A'} \pi_i (\alpha' - \alpha) c_i - \sum_{i \in A \setminus A'} \alpha \pi_i c_i \right] \\
&\quad + \sum_{i \in A \setminus A'} \pi_i [m_{h(i)} v_h T_i - m_{c(i)} v_p (1 - \alpha) c_i] + \sum_{i \in F \cup A'} \pi_i m_{c(i)} v_p (\alpha - \alpha') c_i \\
&= \underbrace{\sum_{i \in F \cup A'} \pi_i v_p (m_{c(i)} - m_p) (\alpha - \alpha') c_i}_{\text{Attention-weighted net value of higher coverage}} \\
&\quad + \underbrace{\sum_{i \in A \setminus A'} \pi_i [m_{h(i)} v_h T_i - m_{c(i)} v_p (1 - \alpha) c_i - m_p v_p \alpha c_i]}_{\text{Attention-weighted net value of higher access}}. \tag{III.11}
\end{aligned}$$

Consider the last equality in (III.11). The first part compares the larger coverage of treatment costs for diseases $i \in F \cup A'$ under health plan 1 to the premium increase necessary to finance this larger coverage. If these differences between the two plans receive full attention $m_{c(i)} = m_p = 1$, they cancel out each other under risk neutrality. If the premium dimension p receives more attention than the co-payment dimensions $c(i)$, this first part of (III.11) is strictly negative. The second part arises if the lower coverage by plan 2 entails a loss in access to treatment. In that case, plan 2 is associated with worse health outcomes in case of sickness as the decision-maker is unable to afford treatment for diseases $i \in A \setminus A'$ when insured under plan 2: a clear disadvantage of the

second plan. Yet, given that there is no treatment, there cannot be any co-payment for these treatments under health plan 2 either: an advantage of plan 2 over plan 1. Finally, as no treatment for diseases $i \in A \setminus A'$ is sought under plan 2, this allows a premium reduction of the entire expected coverage cost $\pi_i \alpha c_i$ compared to plan 1: again an advantage of plan 2. Note that, if α' and α are such that $A = A'$, i.e. if the reduction in coverage does not entail a reduction in access, the second part of (III.11) vanishes since $A \setminus A' = \emptyset$.

Consider the following assumption reminiscent of Assumption III.1:

Assumption III.2.

$$v_p P > \{[\pi_i v_p \alpha c_i]_{i \in F}, [\pi_i v_h T_i]_{i \in A}\}. \quad (\text{III.12})$$

Assumption III.2 states that the disutility from the premium payment for the incumbent plan is larger than each *individual* expected benefit from having (partial) cost coverage and access to treatment. Now we can state the following proposition regarding the possibility for undercutting.

Proposition III.2. Profitable Undercutting

Suppose insurance is voluntary and a single incumbent plan with coverage rate α with $A \neq \emptyset$ is demanded in the absence of any other plan.

(i) If there exists an $\alpha' < \alpha$ such that $A = A'$, and if Assumption III.2 holds for the incumbent plan, a strictly more profitable plan with lower coverage can be constructed that customers will mistakenly choose over the incumbent plan.

(ii) If the incumbent plan features a coverage rate α such that $A' \subset A$ for all $\alpha' < \alpha$, and if Assumption III.2 holds for the incumbent plan, profitable undercutting is possible if $m_{h(i)} = 0$ for the disease(s) $i \in A \setminus A'$.

Proof. See Appendix. □

We call customers' choices described in Proposition III.2 mistaken since they would be better off choosing the high-quality incumbent plan instead of the low-quality plan. The reason for the described undercutting strategy to work is the decision-maker's focus on the premium dimension when making his choice. While the advantage of the new plan over the incumbent plan is concentrated in the premium dimension its disadvantages are distributed across many dimensions. An insurer may thus "hide" the shortcomings of a (qualitatively) disadvantageous insurance plan by reducing coverage rates (or services covered) only slightly for each disease. These service deteriorations in cumulation allow the insurance provider to offer a significant premium reduction. As each single service deterioration is small, customers will not recognize each of them. This allows the firm to retain some of the cost savings from the decrease in quality.

If insurance is voluntary, the attention paid to the health consequences $h(i)$ of insurance choice is independent of the undercutting plan. Thus, when comparing the two plans a customer may pay attention to health consequences of diseases $i \in F \cup A'$ in which the two insurance plans do not differ.²⁴ At the same time, a customer may neglect health consequences $h(i), i \in A \setminus A'$ of diseases in which the two options do actually differ. If access is lost for one of these diseases under the low-quality plan, it will remain unrecognized by the customer. To give an example: if a customer worries particularly about being insured against costs of treatment of common diseases, such as pneumonia, he will particularly look for these features in an insurance plan. The cheaper plan may then be chosen if it covers these common diseases even if it lacks coverage for treatment of some rare types of cancer, and coverage of cancer would, in isolation, be preferred by the customer. Yet, as the customer is so much preoccupied about receiving coverage for the common diseases he neglects to recognize the limited coverage for rare diseases of the cheaper plan.²⁵

Proposition III.2 shows that limited attention may result in a quality deterioration in health insurance markets. Absent of switching costs, a qualitative race to the bottom may arise in markets for complex insurance products. There is a discussion about whether insurance companies “hide” limited coverage in the fine print. This section suggests that it is not font size that makes insurance contracts hard to evaluate, but the sheer size of the contracts. And it is this degree of complexity that allows firms to hide quality reductions in the “fine print”.

Is such undercutting a real danger? After all, the existing literature predominantly finds considerable reluctance to switch between health plans.²⁶ The possibility of undercutting might then not be too serious. Yet, given that this very literature usually calls for interventions to reduce switching costs in order to spur efficiency, the danger of inefficient undercutting absent switching costs should at least be considered. Studies that investigate the reasons of those customers who actually do switch health plans find that the premium plays a suboptimally large role.²⁷ This is striking as there is a considerable number of dimensions in which plans can provide better quality given a premium yet there is only one way to make a health plan cheaper given a level of quality

²⁴This does not mean that attention is “wasted”. The consideration of access is important for the decision-maker to evaluate the desirability of any one of the insurance options against the outside option. If the decision-maker neglected the health consequences of being insured, he would always opt out of insurance as he would disregard all the advantages of being insured.

²⁵Similarly, people may exhibit a tendency to focus on whether their current medication ($\pi_i = 1$) is covered when selecting a Medicare Part D plan. This may result in a failure to consider in addition whether a plan covers those medications these people most likely need in the future ($\pi_i < 1$).

²⁶See e.g. Heiss, Leive, McFadden, and Winter (2012) and Frank and Lamiraud (2008).

²⁷See e.g. Abaluck and Gruber (2011) or Heiss, Leive, McFadden, and Winter (2012).

(i.e. coverage). Thus, calls to decrease switching costs based on efficiency arguments should ascertain that health plan choice absent switching cost indeed optimally weighs price differences against quality differences.²⁸

Both the result of undervaluation of insurance and of the possibility to undercut are the result of an unequal distribution of advantages and disadvantages of one alternative over another across dimensions. As the benefits of insurance are scattered across dimensions, while costs are concentrated in one dimension, the first tend to be underappreciated. Similarly, the undercutting strategy is successful as it concentrates the advantage over a different insurance plan in one dimension (the premium) while spreading the disadvantages across several dimensions. While this section discusses the suboptimal attraction of customers to plans with lower quality and lower premium, we next want to establish that customers with limited attention may even end up buying plans for which lower quality is not even partially compensated by a premium reduction.

4. DOMINATED CHOICES

While most behavioral patterns can be rationalized by some sort of preference, one type of behavior is rather difficult to reconcile with preference-based explanations. If we observe people actively choosing an alternative that fares at most equally well on all dimensions, but is inferior in at least one dimension compared to another available alternative, we remain with two possibilities: the decision-maker does not care at all about the dimension in which the chosen alternative is inferior, or the decision-maker has made a mistake. More precisely, the decision-maker chose a dominated option.

Such dominated choices have been observed in markets for health insurance.²⁹ The model of limited attention proposed here can explain such behavior. For this, assume that plans are described by their premium P and, for each disease $i \in I$, by the degree of coverage $\alpha_i \in [0, 1]$ they offer. The choice set Γ thus comprises different plans g as elements, where each plan g is described by a premium and a vector of coverage rates: $(P, (\alpha_i)_{i \in I})$. If the decision-maker has the option not to insure this can be represented by a “plan” $g^0 \in \Gamma$ with $P = \alpha_i = 0, \forall i \in I$. The following proposition establishes the possibility that a decision-maker may be indifferent between two options, for which one dominates the other.

²⁸Handel (2013) warns against nudging consumers to overcome choice inertia as the improvement in individual choice quality may come at the expense of more severe adverse selection. We complement this by arguing that there might not even be a large benefit of improved individual decision-making to compensate for the exacerbated adverse selection.

²⁹See Sinaiko and Hirth (2011).

Proposition III.3. *Indifference despite Dominance.*

Suppose there exists a choice set Γ of health plans with at least two distinct elements. If there exists a dimension $j \in \{p, c(i) : i \in F, h(i) : i \in \bar{F}\}$ such that $\mu_j > 0$ yet $m_j = 0$. Then there exists a plan that is dominated by one of the available plans, but the decision-maker will be indifferent between these two plan if the dominated plan is added to the choice set.

Proof. See Appendix. □

Let $g^* \in \Gamma$ denote a plan the decision-maker would choose from Γ . Using the idea of Proposition III.3, we can establish the possibility that there exists an option g' that is dominated by g^* , yet would be chosen from the set $\Gamma \cup g'$.

Corollary III.1. *Choice of a Dominated Alternative.*

Suppose there exists a choice set Γ of health plans with at least two distinct elements. If there exists a dimension $j \in \{p, c(i) : i \in F, h(i) : i \in \bar{F}\}$ such that $u(g^*, j) > \min_{g \in \Gamma} u(g, j)$ yet $m_j = 0$. Then there exists a dominated alternative that the decision-maker would choose if it was included in the choice set.

Proof. See Appendix. □

As an illustration, suppose only a single plan is offered that fully covers some set $S \subseteq \bar{F}$ with $|S| \geq 2$, i.e. $\alpha_i = 1, \forall i \in S, \alpha_i = 0, \forall i \in I \setminus S$. Further suppose that the premium is actuarially fair, $P = \sum_S \pi_i c_i$ and the decision-maker is willing to purchase that plan $m_p v_p P \leq \sum_{i \in S} m_{h(i)} v_h T_i$ despite that fact that the decision-maker neglects one of the benefits of coverage: $m_{h(i)} = 0$ for some $i \in S$, say ι . Now, imagine a second plan that is identical to the first plan except for the coverage of disease ι was offered in addition to the first plan, at the same premium P . The introduction of this second plan will not change the attention allocation. Further, the new plan will be considered equally good as the first plan. Since the decision-maker would have chosen the first plan absent the second plan he will now choose either the first or the second plan. He may thus end up choosing the second plan despite it being dominated by the first plan for the simple reason that he happens to neglect exactly the dimension in which the domination occurs.

The corollary and the simple example highlight how the model can naturally explain the observation of dominated choices through modeling neglect. A sufficient condition for dominated choices is stated here since it is obvious, and stated here without proof, that a necessary condition for dominated choices is neglect. For a utility-maximizing decision-maker will only choose a dominated option if he happens to neglect those dimensions that produce the domination. It is important to note the necessity of

neglect for the explanation of dominated choices. Other approaches that also feature biases in the weighting process of different problem dimensions such as Kőszegi and Szeidl (2013) or Bordalo, Gennaioli, and Shleifer (2012) are not capable of explaining such behavior. Though parameter values can be found such that decision weights in these models are approximately zero, they can never be exactly zero. Yet, this is necessary to model neglect and dominated choices as one of its behavioral consequence.

Note that the model allows for stronger predictions than stating the mere possibility of dominated choices. The requirement that some benefits of the dominating insurance plan are neglected allows to make further predictions. First, the model predicts at most indifference between a dominating and a dominated alternative. It cannot happen that the decision-maker strictly prefers the dominated over the dominating alternative. This seems plausible: limited attention may attenuate to which extent an advantage is appreciated. It should not lead to an advantage being misperceived as a disadvantage. Second, the model predicts dominated choices never to occur in binary choices. In binary choices, at least one dimension in which the dominated alternative is inferior must be considered. This suffices for a dominated alternative never to be chosen. Finally, if a dominated alternative is chosen, this alternative and the dominating alternative must share advantages over a third alternative that distract attention from the dimensions in which the domination occurs. This again implies that a dominated alternative that is weakly worse than all available alternatives, i.e. an alternative dominated by all other alternatives, will never be chosen. It also means that the shared advantages over the third alternative must be large enough compared to the disadvantage(s) of the dominated alternative, for otherwise the first could not distract from the latter. In this sense, the disadvantages of the dominated alternative that is chosen have to be minor.

5. SPECIFIC VS. COMPREHENSIVE INSURANCE: THE BENEFITS OF ISOLATING RISKS

It has been observed that people's willingness-to-pay (WTP) for various specific insurances exceeds their willingness-to-pay for a comprehensive insurance that covers all of the incidences the specific insurances are addressing. Johnson, Hershey, Meszaros, and Kunreuther (1993) provide several examples of such an unbundling effect.³⁰ They ar-

³⁰In a series of experiments they ask their subjects for their willingness-to-pay for health insurance that covers hospitalization either for any disease, for any accident, for any reason, or for any disease or accident. They find that if subjects are asked their WTP for any disease (followed by any accident), or asked their WTP for any accident (followed by any disease), the sum of these two WTP (\$89.10 and \$69.55 on average) significantly exceeds the WTP expressed for the insurance covering any reason (\$41.53 on average) or the insurance covering any disease or accident (\$47.12

gue that this effect is due to a greater availability of the more specific events compared to the unspecific “any reason”. Our model may complement the availability hypothesis with an explanation based on complexity-reduction.

Suppose that people tend to think of three categories of consequences in which the alternatives differ when considering this decision problem: premium, coverage in case of accident, and coverage in case of disease.³¹ Denote by P the premium, c_a the cost treatment in case of an accident, π_a the probability the DM associates with having an accident, c_d the cost of treatment in case of a disease, and π_d the probability to be hospitalized for a disease.

First, consider the case of an insurance that only covers one of the incidences (either accident or disease). The decision problem of whether to buy such an insurance comprises two dimensions: the premium dimension and the dimension associated with the payment in case of disease. Each option has an advantage in exactly one dimension. If the decisions are made sequentially without prior anticipation of the second decision³², the decision-maker solves two two-dimensional problems with advantages and disadvantages being condensed in one dimension each. Yet, if the decision-maker has to choose between insuring against any disease or accident, he solves one three-dimensional problem. The advantages of being insured are spread across two dimension: payment in case of an accident and payment in case of a disease. The disadvantage is condensed in only one dimension: the premium. The attention process then favors remaining uninsured. The reason for this effect is that the integration of several incidences into one comprehensive insurance makes the insurance decision more complex. In particular, it adds a benefit dimension while integrating the cost into an already existing dimension (the premium). As a decision-maker cannot fully take into account all dimensions he concentrates on those with the largest utility differences. This tends to be the premium dimension as it integrates the costs of covering several incidences. As a result the benefit dimensions are not fully considered.

This unbundling effect is not only present when the two more specialized insurances incorporate only a single benefit dimension. We can show that any “split” of a comprehensive insurance into an arbitrary number of more specialized insurance plans will result in an increased willingness-to-pay.

on average).

³¹For simplicity, we assume treatment cost to be affordable for both incidences. The argument can easily be replicated for the cases in which treatment for one or both incidents is unaffordable without insurance.

³²Alternatively, one may assume that the decision-maker narrowly brackets such that he solves the two choice problems in isolation.

Proposition III.4. *Unbundling of Insurance Plans*

Let $S \subseteq I$, $|S| \geq 2$ be the set of diseases for which a comprehensive insurance plan offers full coverage of treatment cost. Let (S_1, \dots, S_z) be a partition of S and let $W(S)$ be the maximum willingness-to-pay for an insurance plan covering the treatment costs for all diseases $i \in S$. Then $W(S) < \sum_{i=1}^z W(S_i)$.

Proof. See Appendix. □

Unbundling a comprehensive insurance mitigates the extent of underappreciation of value. This underappreciation of comprehensive insurance might pose a dilemma to insurance providers. On the one hand, there is an incentive to split insurance plans into more specific plans in order to mitigate the underappreciation of the value of insurance. On the other hand, customers may be reluctant to consider a large number of specific plans individually.

6. INSURANCE PROVISION IN MARKET EQUILIBRIUM

We want to investigate which insurance plans are offered in equilibrium in a regulated market. Due to the discontinuities in the attention function m_j , that result in discontinuities in the payoff functions of firms, we cannot be certain that equilibria always exist. We therefore look at a particular setting. Each firm can offer only a single plan. We confine attention to the case in which firms choose the diseases for which they provide coverage. Yet, if they choose to cover a particular disease, they are bound to cover the full treatment cost.³³ A firm's plan choice is then a set $S \subseteq I$ of diseases covered and a premium P . This simplification allows the convenience to define the benefit b_i of having insurance for disease $i \in I$ by

$$b_i = \begin{cases} \pi_i v_p c_i & \text{if } i \in F, \\ \pi_i v_h T_i & \text{if } i \in \bar{F}. \end{cases}$$

Since there are no differences between having or not having insurance in the monetary consequences for diseases $i \in \bar{F}$ and no differences between having and not having insurance in the health consequences for diseases $i \in F$, we will write m_i to denote the attention paid to the benefit of having insurance for disease i , i.e. $m_i = m_{c(i)}$, $\forall i \in F$ and $m_i = m_{h(i)}$, $\forall i \in \bar{F}$.³⁴

³³One could think of this as a regulatory requirement to eliminate the undercutting incentives we discussed in Proposition III.2.

³⁴We assume throughout the section that the premium payments do not result in an inability to afford treatments that would be affordable without insurance. Formally, with the decision-maker's budget being B we assume $B - P(S) > c_i$, $\forall F \setminus S$. Since we will consider actuarially fair premia in this

We assume that customers are equal with regard to their preferences, their risk, and their cognitive abilities. These customers choose the plan that maximizes their decision utility given the choice set they face. If more than one plan maximizes decision utility, demand is split equally among the maximizing plans unless noted otherwise.

Suppose insurance is voluntary, i.e. there exists the outside option not to insure at all $(S, P) = (\emptyset, 0)$. Consider the following equilibrium candidate. For each subset $S \subseteq I$ a plan covering this very subset is offered by more than one firm at an actuarially fair premium, i.e. $P = \sum_{i \in S} \pi_i c_i$. Then, if $\sum_{i \in I} \pi_i c_i > b_i \forall i \in \bar{F}$ is satisfied, this constitutes an equilibrium.

Proposition III.5. *If $\sum_{i \in I} \pi_i c_i > b_i \forall i \in \bar{F}$, then there always exists an asymmetric equilibrium in which each possible plan with an actuarially fair premium $(S, P) : S \subseteq I, P = \sum_{i \in S} \pi_i c_i$ is offered by at least two firms.*

Proof. See Appendix. □

It is interesting to investigate which plans are purchased by customers in the above equilibrium. It turns out that equilibrium insurance, and thus the welfare properties of the equilibrium, strongly depend on the structure of insurance benefits b_i . We will assume that $b_i \neq b_{i'} \forall i, i' \in I : i \neq i'$.

The following proposition characterizes the insurance plan that is purchased in equilibrium.³⁵

Proposition III.6. *In the equilibrium described in Proposition III.5 the customers purchase the plan (S^*, P^*) with*

$$S^* = \{i \in \bar{F} : b_i - v_p \pi_i c_i \geq \kappa_{r(i)}\}$$

$$\text{with } r(i) = |\{i' \in I : b_{i'} > b_i\}| + 1,$$

$$P^* = \sum_{i \in S^*} \pi_i c_i$$

Proof. See Appendix. □

Note the following properties of insurance purchase in equilibrium. First, in equilibrium no diseases are insured for which customers do not need insurance: $S^* \subseteq \bar{F}$. This

section this translates into $B - \sum_S \pi_i c_i > c_i, \forall F \setminus S$. Again, this abstracts from any issues related to the affordability of premia. We note that this is not without loss of generality for a treatment may become unaffordable because insurance for other diseases cuts deep into the decision-maker's budget. We maintain this assumption in order to focus on the implications of limited attention on equilibrium insurance.

³⁵We assume that if a customer is indifferent between a plan $(S, \sum_S \pi_i c_i)$ and a plan $(S', \sum_{S'} \pi_i c_i)$ with $S' \subset S$, then the customer purchases the more comprehensive plan covering S .

is true because $b_i - v_p \pi_i c_i - \kappa_{r(i)} = -\kappa_{r(i)} < 0$ holds for all $i \in F$. However, S^* might be the empty set, i.e. customers may not insure at all in equilibrium. This means it is possible that none of the gains from trade are realized in equilibrium. This happens if there does not exist $i \in \bar{F}$ satisfying the condition for inclusion in S^* . Such a situation occurs if the largest insurance benefits are produced by diseases with affordable treatments. In that case, diseases for which insurance is necessary, $i \in \bar{F}$, receive very low attention ranks and therefore very high attention thresholds $\kappa_{r(i)}$ that may preclude them from being appreciated in the decision process.³⁶

The comprehensive plan plays an important role in the equilibrium as it fixes the attention allocation and makes it distraction-proof. This stabilization of attention comes at a cost however. Limited customer attention may be wasted on diseases for which customers do not need (and do not buy) insurance. Formally, there may be diseases $i \in F$ with $m_i > 0$, while we know that $S^* \cap F = \emptyset$. A straightforward question would be whether there always exist equilibria in which the market offers coverage only for diseases $i \in \bar{F}$. Such a setting would be desirable as attention would only be allocated to diseases for which insurance is beneficial. Second, under such an attention allocation the extent of coverage $S^* \subseteq \bar{F}$ that customers eventually choose could be larger since they do not waste attention considering coverage for unnecessary insurance. Unfortunately, this is not the case. It can be shown that in such an environment profitable deviations may occur in which firms distract attention by offering and selling plans that include unnecessary coverage while excluding necessary coverage.³⁷ In that sense, by stabilizing attention the comprehensive option ensures quality in the market. Ironically, it does so by offering the highest extent of unnecessary coverage, i.e. it includes coverage for all $i \in F$. We want to underline that the comprehensive plan may ensure quality in the market without ever being chosen itself.

Let us consider the welfare properties of the equilibrium. To do so, let us first note that welfare is maximized by any set $S \in I$ that maximizes $\sum_{i \in S} (b_i - \pi_i c_i)$. This means that the first-best is characterized by

$$S^{fb} \supseteq \bar{F}. \quad (\text{III.13})$$

That means, customers would purchase insurance for all diseases for which they need insurance, $i \in \bar{F}$. Comparing this to the insurance plan purchased in the equilibrium

³⁶It is interesting to note that the three plans that support this equilibrium are the outside option $(\emptyset, 0)$, the comprehensive plan $(I, \sum_i \pi_i c_i)$, and the chosen plan (S^*, P^*) . The former two fix the attention allocation and ensure that attention cannot be distracted. The latter is the one to which customers assign highest decision utility given that attention allocation and it may happen that it coincides with one of the former.

³⁷A proof can be found in the Appendix.

with limited attention shows that $S^* \subseteq \bar{F} \subseteq S^{fb}$. While competition drives premia to actuarially fair levels, customers tend to underinsure under limited attention. In the most extreme case, customers may choose not to insure at all, $S^* = \emptyset$, thereby foregoing all the benefits from trade prevalent in this market: $\sum_{\bar{F}}(b_i - \pi_i c_i)$.

We may wonder whether the equilibrium is at least constrained-efficient in the sense that it maximizes welfare given the cognitive limitations of the market participants. Such a constrained-efficient outcome would result in an insurance of a set $S \subseteq I$ that maximizes $\sum_{i \in S}(b_i - \pi_i c_i)$ under the constraint that insurance is purchased voluntarily: $\tilde{U} \geq 0$. Unfortunately, the equilibrium set S^* generally does not coincide with the constrained-efficient set S^C either.³⁸ One reason for this is that in the competitive equilibrium attention is wasted towards considering diseases i for which insurance is not necessary. In addition, among the diseases $i \in \bar{F}$, attention is attracted towards diseases with large benefits b_i . These do not necessarily coincide with the diseases that deliver the largest welfare gains of insurance $b_i - \pi_i c_i$. Finally, S^* maximizes decision-utility, while S^c maximizes experienced utility (while keeping decision utility nonnegative). Thus S^c may include treatments that result in an increase in experienced utility, yet a decline in decision-utility, as long as this decline does not make the whole insurance purchase undesirable.

7. CONCLUSION

This chapter seeks to illustrate how four phenomena that have been observed in the choice of health insurance may be the result of the complexity inherent in this choice problem. First, we have shown that people with a tendency to simplify complex decisions tend to underappreciate the value of health insurance. Second, this tendency to simplify allows firms to “hide” quality reductions by scattering them across many attributes of the health plan. Third, their propensity to neglect may lead people to make dominated choices. Finally, we have indicated an incentive to unbundle comprehensive health plans in order to mitigate the extent of underappreciation of value.

These results may give rise to a dilemma faced by policy makers who seek to increase insurance coverage. If one acknowledges an underappreciation of value one may support calls for an individual mandate for health insurance incorporating a couple of minimum quality standards. However, to ensure a first-best allocation an individual mandate would have to be coupled with mandated benefits for *all* diseases for which individuals cannot self-insure. Insurance providers would only be able to compete on the premium

³⁸To show this consider that S^* can be empty, while the constrained-efficient set cannot. For any $i \in \bar{F}$, the insurance plan $(i, \pi_i c_i)$ results in strictly positive decision utility absent any alternative insurance plan (save the outside option).

dimension. Such a policy has its own drawbacks though. Such a restrictive policy would forgo the benefits of product differentiation for different tastes/risks. In addition, it would create a strong incentive for health providers to lobby for their products to be covered by mandated benefits. Finally, the model of limited attention identifies the underappreciation of the value of such a mandated comprehensive insurance as a particular obstacle. If people underappreciate the benefits while focusing on the cost, such a policy will be highly unpopular.

In sum, acknowledging the existence of the described biases in choice behavior may support calls for policy interventions mandating insurance with extensive benefits. Yet, the very existence of these biases will make such policies quite unpopular. Limited attention as modeled here may thus not only drive a wedge between the need for and the acceptance of policy interventions but even make them reciprocal.

This work is an attempt to model the complexity involved in choosing a health plan that goes beyond modeling the number of choices as the main source of complexity. We argue that a major part of the complexity involved is due to the many aspects this choice problem has. Given that we only consider the large number of diseases a health plan may or may not cover, and that health plans may vary on many more attributes, we are confident that further research into this aspect of complexity will be instructive. Finally, we think that research on limited attention may add nicely to the well-researched topic of adverse selection in insurance markets. If people underappreciate insurance they may select out of a market to a degree that is suboptimal both from a social and an individual point of view. In addition, selection may occur both due to a heterogeneity in underlying risks and a heterogeneity in cognitive capabilities. We conclude that further research into the issue of complexity in the market for health insurance can take several promising directions.

A. Appendices

1. APPENDIX TO CHAPTER I

Equivalence of Two Approaches to Calculate Gain-Loss Utility

KR derive a percentile-wise comparison of material outcomes for some belief $F(m)$ over outcomes. Formally, they define the material outcome at percentile π :

For any distribution F over \mathbb{R} and any $\pi \in (0, 1)$ let $m_F(\pi)$ be the material utility level at percentile π , defined implicitly by

$$\begin{aligned} (i) & F(m_F(\pi)) \geq \pi, \\ (ii) & F(m) < \pi \text{ for all } m < m_F(\pi). \end{aligned}$$

This definition yields a unique function $m_F(\pi) : (0, 1) \rightarrow \mathbb{R}$ with

$$m_F = \begin{cases} m_1 & \forall \pi \in (0, F(m_1)] \\ m_2 & \forall \pi \in (F(m_1), F(m_2)] \\ \dots & \\ m_k & \forall \pi \in (F(m_{k-1}), 1) \end{cases} \quad (\text{A.1})$$

where (m_1, \dots, m_k) constitutes the support of $F(m)$ with $m_1 < m_2 < \dots < m_k$.

Proof. First, $m_F = m_1$, $\forall \pi \in (0, F(m_1)]$. Suppose not. In this case, there exists a $\pi \in (0, F(m_1)]$ for which either $m_F(\pi) < m_1$ or $m_F(\pi) > m_1$. If $m_F(\pi) < m_1$ there is a contradiction to (i), since $F(m_F(\pi)) = 0 < \pi$, $\forall \pi \in (0, F(m_1)]$. If $m_F(\pi) > m_1$ there is a contradiction to (ii), since there exists a level of $m < m_F(\pi)$, namely m_1 , for which $F(m) \geq \pi$.

Similarly, for all m_i , $i = 2..k - 1$, $m_F = m_i$, $\forall \pi \in (F(m_{i-1}), F(m_i)]$. Suppose not. Then there exists a $\pi \in (F(m_{i-1}), F(m_i)]$ such that either $m_F(\pi) < m_i$ or $m_F(\pi) > m_i$. If $m_F(\pi) < m_i$ there is a contradiction to (i) as $F(m_F(\pi)) < \pi$. If $m_F(\pi) > m_i$ there is a contradiction to (ii), since there exists an $m < m_F(\pi)$, namely m_i such that $F(m) \geq \pi$.

Finally, $m_F(\pi) = m_k$, $\forall \pi \in (F(m_{k-1}), 1)$. Suppose not. Then there exists a $\pi \in$

$(F(m_{k-1}, 1))$ such that either $m_F(\pi) < m_k$ or $m_F(\pi) > m_k$. The first case contradicts (i) since $F(m_F(\pi)) < \pi$. The second case contradicts (ii) since there exists an $m < m_F(\pi)$, namely m_k such that $F(m) \geq \pi$. \square

It can thus be shown that the function m_F , implicitly defined by KR, can explicitly be defined as the *quantile function* of the cumulative distribution function F .

KR define gain-loss utility in the following way: at each percentile the material outcome of the new belief is compared to the material outcome of the old belief. If the former is higher than the latter the individual experiences a gain at this percentile, otherwise it experiences a loss at this percentile. Formally, define psychological utility from a change in belief by

$$\tilde{v}(F_t, F_{t-1}) = \eta \int_0^1 \mu [m_{F_t}(\pi) - m_{F_{t-1}}(\pi)] d\pi$$

with:

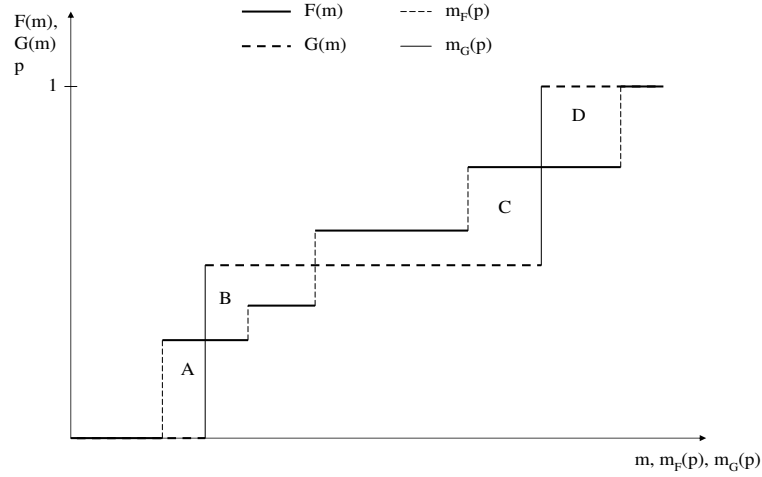
$$\begin{aligned} \mu(z) &= z \text{ if } z \geq 0, \\ \mu(z) &= \lambda z \text{ if } z < 0. \end{aligned}$$

Take as an illustration Figure A.1 and suppose F represents a posterior (F_t) and G the prior belief (F_{t-1}). Then $\tilde{v}(F, G) = -\lambda A + B - \lambda C + D$.

It is straightforward that the same result can be obtained by comparing the old and new belief directly without deriving the quantile functions as long as μ is two-piece linear as assumed. One can calculate gain-loss utility instead by assigning a gain to any level of material outcome to which the new belief, say F , assigns lower density than the old belief (G), i.e. when $F(m) < G(m)$. A loss is assigned to any level of material outcome to which the new belief (F) assigns higher density than the old belief (G), i.e. when $F(m) > G(m)$. Thus, one can calculate the gain-loss utility directly by comparing the beliefs at each level of material outcome.

$$v(F_t, F_{t-1}) = \eta \int_{-\infty}^{\infty} \mu [F_{t-1}(m) - F_t(m)] dm.$$

With above definition, $v(F, G) = -\lambda A + B - \lambda C + D$. Thus, both approaches yield the same result as they both calculate equally-sized losses for the areas A and C and equally-sized gains for the areas B and D. The second approach, however, spares the detour over deriving the quantile functions $m_F(\pi)$, $m_G(\pi)$ first. Note, however, that this simplification is only valid when assuming the two-piece linear representation of gain-loss utility μ .

Figure A.1.: Comparing F and G versus comparing m_F and m_G .

Proof of Lemma I.1

Proof. To prove Lemma 1, one needs to show that (i) there exist beliefs p_1 for which $\mathbb{E}_1 U^1(NT|NT) - \mathbb{E}_1 U^1(T|NT)$, $V_{NT} = \mathbb{E}_1 U^1(NT|NT) - \mathbb{E}_1 U^1(T|T)$, and $\mathbb{E}_1 U^1(T|T) - \mathbb{E}_1 U^1(NT|T)$ equal zero respectively, (ii) these beliefs are unique, and (iii) have the order described by Lemma 1. Define

$$\Gamma_{NT} \equiv \mathbb{E}_1 U^1(NT|NT) - \mathbb{E}_1 U^1(T|NT), \quad (\text{A.2})$$

$$\Gamma_T \equiv \mathbb{E}_1 U^1(T|T) - \mathbb{E}_1 U^1(NT|T). \quad (\text{A.3})$$

Furthermore, recall the definition $V_{NT} \equiv \mathbb{E}_1 U^1(NT | NT) - \mathbb{E}_1 U^1(T | T)$.

(i) Existence

The proof is simple. Note that $\Gamma_{NT} > 0$, $V_{NT} > 0$ and $\Gamma_T < 0$ for $p_1 = 0$ while $\Gamma_{NT} < 0$, $V_{NT} < 0$, and $\Gamma_T > 0$ for $p_1 = 1$. By continuity of Γ_{NT} , V_{NT} and Γ_T , the functions have at least one root on the interval $(0, 1)$. Denote by p_{NT}^*, p^*, p_T^* a root of $\Gamma_{NT}, V_{NT}, \Gamma_T$ respectively.

(ii) Uniqueness

It can be shown that the first derivatives of Γ_{NT} and V_{NT} with respect to p_1 are strictly negative at their respective root while the first derivative of Γ_T with respect to p_1 is strictly positive at its root. This, coupled with continuity, implies the uniqueness of the roots.

(ii,a) Γ_{NT} is strictly decreasing in p_1 at $\Gamma_{NT} = 0$.

At $p_1 = p_{NT}^*$, $\Gamma_{NT} = \mathbb{E}_1 U^1(NT|NT) - \mathbb{E}_1 U^1(T|NT) = 0$. This implies

$$\begin{aligned} (1 - p_{NT}^*)\Delta_h - p_{NT}^*\Delta_s - p_{NT}^*(1 - p_{NT}^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) \\ + \gamma_1\eta[\lambda(1 - p_{NT}^*)\Delta_h - p_{NT}^*\Delta_s] &= 0 \\ \Rightarrow \chi_1 \equiv (\Delta_h + \Delta_s) - p_{NT}^*\eta(\lambda - 1)(\Delta_h + \Delta_s) + \gamma_1\eta(\lambda\Delta_h + \Delta_s) \\ &= \frac{1}{1 - p_{NT}^*}\Delta_s(1 + \gamma_1\eta) > 0. \end{aligned}$$

Differentiating Γ_{NT} w.r.t. p_1 yields

$$\frac{\partial\Gamma_{NT}}{\partial p_1} = (\Delta_h + \Delta_s) - (1 - 2p_1)\eta(\lambda - 1)(\Delta_h + \Delta_s) + \gamma_1\eta(\lambda\Delta_h + \Delta_s).$$

At $p_1 = p_{NT}^*$ this differential is negative as

$$\begin{aligned} \frac{\partial\Gamma_{NT}(p_{NT}^*)}{\partial p_1} &= -(\Delta_h + \Delta_s) - (1 - 2p_{NT}^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) - \gamma_1\eta(\lambda\Delta_h + \Delta_s) \\ &= -\chi_1 - (1 - p_{NT}^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) < 0. \end{aligned}$$

(ii,b) V_{NT} is decreasing in p_1 at p^* .

At p^* , $V_{NT} = 0$. This implies

$$\begin{aligned} V_{NT} = -\Delta_s + (1 - p^*)[(\Delta_h + \Delta_s) - p^*\eta(\lambda - 1)(\Delta_h + \Delta_s)] &= 0 \\ \Rightarrow \chi_2 \equiv (\Delta_h + \Delta_s) - p^*\eta(\lambda - 1)(\Delta_h + \Delta_s) &= \frac{\Delta_s}{1 - p^*} > 0. \end{aligned}$$

Differentiating V_{NT} w.r.t. p_1 yields

$$\frac{\partial V_{NT}}{\partial p_1} = -(\Delta_h + \Delta_s) - (1 - 2p_1)\eta(\lambda - 1)(\Delta_h + \Delta_s).$$

At $p_1 = p^*$ this differential is negative as

$$\begin{aligned} \frac{\partial V_{NT}(p^*)}{\partial p_1} &= -(\Delta_h + \Delta_s) - (1 - 2p^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) \\ &= -\chi_2 - (1 - p^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) < 0. \end{aligned}$$

(ii,c) Γ_T is increasing in p_1 at p_T^* .

At p_T^* , $\Gamma_T = 0$. This implies

$$\begin{aligned} \Gamma_T = p_T^*\Delta_s - (1 - p_T^*)\Delta_h + p_T^*(1 - p_T^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) \\ + \gamma_1\eta[-(1 - p_T^*)\Delta_h + \lambda p_T^*\Delta_2] &= 0 \\ \Rightarrow \chi_3 \equiv (\Delta_h + \Delta_s) - p_T^*\eta(\lambda - 1)(\Delta_h + \Delta_s) + \gamma_1\eta(\Delta_h + \lambda\Delta_s) \\ &= \frac{1}{1 - p_T^*}\Delta_s(1 + \gamma_1\eta\lambda) > 0. \end{aligned}$$

Differentiating Γ_T w.r.t. p_1 yields

$$\frac{\partial \Gamma_T}{\partial p_1} = (\Delta_h + \Delta_s) + (1 - 2p_1)\eta(\lambda - 1)(\Delta_h + \Delta_s) + \gamma_1\eta(\Delta_h + \lambda\Delta_s).$$

At $p_1 = p_T^*$ this differential is positive as

$$\begin{aligned} \frac{\partial \Gamma_T(p_T^*)}{\partial p_1} &= (\Delta_1 + \Delta_2) + (1 - 2p_T^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) + \gamma_1\eta(\lambda\Delta_h + \Delta_s) \\ &= \chi_3 + (1 - p_T^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) > 0. \end{aligned}$$

(iii) Order: $0 < p_T^* < p^* < p_{NT}^* < 1$.

To complete the proof of the Lemma, it needs to be shown that the roots have the stated order. Note that $\Gamma_{NT} > 0$ and $\Gamma_T < 0$ for $p_1 = 0$ and $\Gamma_{NT} < 0$ and $\Gamma_T > 0$ for $p_1 = 1$. This implies that the roots are bounded away from zero and one. Furthermore, it can be shown that both $\Gamma_{NT} > 0$ and $\Gamma_T > 0$ at $p_1 = p^*$. These observations, in addition to the existence of unique roots and continuity of the functions, imply that $p_T^* < p^*$ and $p_{NT}^* > p^*$.

(iii,a) $\Gamma_{NT}(p^*) > 0$.

$$\begin{aligned} \Gamma_{NT}(p^*) &= V_{NT}(p^*) - \gamma_1\eta[-\lambda(1 - p^*)\Delta_h + p^*\Delta_s] \\ &= 0 + \gamma_1\eta[\lambda(1 - p^*)\Delta_h - p^*\Delta_s] \end{aligned}$$

Note that $r_{NT} - r_T = \Delta_h + \Delta_s > 0$. This means that at p^* , $V_{NT} = 0$ implies

$$\begin{aligned} V_{NT} &= (1 - p^*)\Delta_h - p^*\Delta_s - p^*(1 - p^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) = 0 \\ &\Leftrightarrow (1 - p^*)\Delta_h - p^*\Delta_s > 0 \\ &\Rightarrow \lambda(1 - p^*)\Delta_h - p^*\Delta_s > 0. \end{aligned}$$

Thus Γ_{NT} is positive at p^* .

(iii,b) $\Gamma_T(p^*) > 0$

At $p_1 = p^*$, $\Gamma_T(p^*) = \gamma_1\eta[-(1 - p^*)\Delta_h + \lambda p^*\Delta_s]$. For this to hold it is necessary that $(1 - p^*)\Delta_h < \lambda p^*\Delta_s$. Note that, by definition of p^* , $(1 - p^*)\Delta_h = p^*\Delta_s + p^*(1 - p^*)\eta(\lambda - 1)(\Delta_h + \Delta_s)$. Using this equation and solving for λ gives

$$\begin{aligned} (1 - p^*)\Delta_h &< \lambda p^*\Delta_s \\ \Leftrightarrow p^*\Delta_s + p^*(1 - p^*)\eta(\lambda - 1)(\Delta_h + \Delta_s) &< \lambda p^*\Delta_s \\ \Leftrightarrow p^*\Delta_s - p^*(1 - p^*)\eta(\Delta_h + \Delta_s) &< \lambda[p^*\Delta_s - p^*(1 - p^*)\eta(\Delta_h + \Delta_s)] \\ &\Leftrightarrow 1 < \lambda \end{aligned}$$

which is true by assumption. □

Derivation of the value of information W

We derive the value $W \equiv \mathbb{E}_0 U^0(i|i) - \mathbb{E}_0 U^0(n|n)$ assuming a prior $p_0 < p^*$ such that NT is the default action. The derivation for the alternative default action proceeds analogously.

$$\begin{aligned} W(NT^d) &= \mathbb{E}_0 [u_2|i] \\ &\quad - \mathbb{E}_0 [u_2|n] \\ &\quad + \mathbb{E}_0 [u_1|i] \end{aligned}$$

$$\begin{aligned} W(NT^d) &= q^- [(1 - p^-)A + p^- B - p^-(1 - p^-)\eta(\lambda - 1)r_{NT}] \\ &\quad + q^+ [(1 - p^+)C + p^+ D - p^+(1 - p^+)\eta(\lambda - 1)r_T] \\ &\quad - [(1 - p_0)A + p_0 B - p_0(1 - p_0)\eta(\lambda - 1)r_{NT}] \\ &\quad - q^- q^+ \gamma_1 \eta(\lambda - 1) [(1 - p^+)(\Delta_h + \Delta_s) + (2p^- - 1)\Delta_s + (p^+ - p^-)r_{NT}] \end{aligned}$$

Noting that, under Bayesian updating, the prior is a convex combination of the potential posteriors, one can replace the third line with

$$\begin{aligned} [(1 - p_0)A + p_0 B - p_0(1 - p_0)\eta(\lambda - 1)r_{NT}] &= \\ [q^-(1 - p^-) + q^+(1 - p^+)] A + [q^- p^- + q^+ p^+] B \\ - [q^- p^- (1 - p^-) + q^+ p^+ (1 - p^+) + q^- q^+ (p^+ - p^-)^2] \eta(\lambda - 1) r_{NT} \end{aligned}$$

This yields

$$\begin{aligned} W(NT^d) &= q^+ [p^+ \Delta_s - (1 - p^+) \Delta_h + p^+(1 - p^+)\eta(\lambda - 1)(r_{NT} - r_T)] \\ &\quad + q^- q^+ \eta(\lambda - 1) r_{NT} (p^+ - p^-)^2 \\ &\quad - q^- q^+ \gamma_1 \eta(\lambda - 1) [(1 - p^+)(\Delta_h + \Delta_s) + (2p^- - 1)\Delta_s + (p^+ - p^-)r_{NT}]. \end{aligned}$$

Using the notation we introduced earlier, we can rewrite this as

$$W(NT^d) = q^+ V_T(p^+) + VoNI(NT^d) + ED(NT^d).$$

Proof of Proposition I.2

Proof. First, by Corollary I.1, perfect information must be instrumental. Second, by Lemma I.2, the default action is NT if $p_0 < p^*$ and T if $p_0 > p^*$. For ignorance to be a PPE, (a) the value of perfect information must be negative given that the plan is not to test: $VoII(\cdot, n) < 0$, and (b) the plan prescribing not to test must be preferred to the plan to test: $W < 0$.

If $p_0 < p^*$, the default action is not to treat (NT). The value of perfect information is then

$$VoPI(NT^d, n) = p_0 \Delta_s [1 + (1 - p_0) \gamma_1 \eta (\lambda - 1) + \gamma_0 \eta] + p_0 (1 - p_0) \eta (\lambda - 1) (1 - \gamma_1) r_{NT}$$

which is strictly positive. Thus ignorance cannot be a PPE for someone with default NT as condition (a) is never satisfied.

Now suppose $p_0 > p^*$, the default action is to treat (T). It is easy to see from equations (I.13) and (I.15) that for perfect information $W < VoII(\cdot, n)$, thus condition (a) implies condition (b). In other words, if rejection of perfect information is a PE then it must be the PPE. The value of perfect information for someone with default T is given by

$$VoPI(T^d, n) = (1 - p_0) \Delta_h (1 - \gamma_0 \eta) + p_0 (1 - p_0) \eta (\lambda - 1) [r_T - \gamma_1 (r_T + \Delta_h)].$$

This value is negative if and only if $\gamma_1 > \frac{r_T}{r_T + \Delta_h}$ and the degree of loss aversion exceeds a critical value λ^* given by

$$\lambda^* \equiv \frac{\Delta_h (1 + \gamma_0 \eta)}{\eta p_0 [\gamma_1 (r_T + \Delta_h) - r_T]} + 1. \tag{A.4}$$

□

Proof of Proposition I.3

Proof. We look at the value of a particular test: a test with $\epsilon^- = 0$ as assumed and $\epsilon^+ = \frac{p_0}{1-p_0} \frac{\Delta_s}{\Delta_h}$. Such a test produces the posterior $p^+ = \Delta_h / (\Delta_h + \Delta_s)$. It is of particular interest as it is the worst test in terms of ϵ^+ (among those with $\epsilon^- = 0$) that is still instrumental from the point of view of the physician. Any test with a smaller false positive rate would be recommended by the physician as it yields a strictly positive value of information from his point of view. We derive a condition under which this particular test is rejected by a patient with reference-dependent preferences. If the value of information for the patient turns out to be negative for such a test, by continuity of $VoII$ it will still be negative for a test with slightly smaller ϵ^+ to which the physician attributes a strictly positive value. Then there exists a range of tests in terms of ϵ^+ to which the physician assigns positive value but the patient assigns negative value.

To derive whether the patient agrees to being tested or refuses to being tested we first need to find out whether the patient would regard the test as instrumental or noninstrumental. The patient will regard the test as instrumental if

$$W(NT^d) - VoNI(NT^d) \geq 0$$

It can easily be verified that this is the case with $p^- = 0$ and $p^+ = \frac{\Delta_h}{\Delta_h + \Delta_s}$. In order for ignorance to be a PPE, it suffices to check whether $W(NT^d) < 0$ as for above posteriors $VoII(NT^d, n) < W(NT^d)$. This means that if $W(NT^d) < 0$ ignorance is both a personal equilibrium and, in addition, the preferred personal equilibrium.

With $p^- = 0$ and $p^+ = \frac{\Delta_h}{\Delta_h + \Delta_s}$, $W(NT^d)$ equals

$$q^+ p^+ (1 - p^+) \eta (\lambda - 1) (\Delta_h + \Delta_s) + VoNI(NT^d)$$

Thus if $VoNI(NT^d)$ is sufficiently negative ignorance is optimal. This is the case if

$$\Delta_s < q^- r_{NT} \left(\gamma_1 - \frac{\Delta_h}{\Delta_h + \Delta_s} \right). \quad (\text{A.5})$$

If condition A.5 is satisfied a patient with default NT refuses a test that results in posterior probabilities of being sick of 0 or $\Delta_h / (\Delta_h + \Delta_s)$. Together with the fact that this patient would not refuse a perfect test yields the result that there exists a threshold $p^+ > \Delta_h / (\Delta_h + \Delta_s)$ below which the patient refuses to being tested. Thus for any prior $p_0 \in \left(0, \frac{\Delta_h}{\Delta_h + \Delta_s} \right)$ for which above condition is satisfied, there exists an interval of false-positive rates ϵ^+ for which the physician recommends the test but the patient refuses. \square

Derivation of Observation 2

An increase in the effectiveness of treatment (a) increases Δ_s , the net benefit of treatment, and (b) decreases r_T , which is the smaller the more effective the treatment. It is possible to capture this effect by differentiating with respect to r_T while assuming $\partial \Delta_s / \partial r_T = -1$.

Proof.

$$\frac{\partial VoNI(NT^d)}{\partial r_T} = 0. \quad (\text{A.6})$$

$$\frac{\partial VoNI(T^d)}{\partial r_T} = q^- q^+ \eta (\lambda - 1) (p^+ - p^-) [(p^+ - p^-) - \gamma_1]. \quad (\text{A.7})$$

$$\left(\frac{\partial W(NT^d)}{\partial r_T} \right)_{\epsilon^+ = \epsilon^- = 0} = -p_0 (1 + (1 - p_0) \gamma_1) < 0. \quad (\text{A.8})$$

$$\left(\frac{\partial W(T^d)}{\partial r_T} \right)_{\epsilon^+ = \epsilon^- = 0} = p_0 (1 - p_0) \eta (\lambda - 1) (1 - \gamma_1) > 0. \quad (\text{A.9})$$

\square

Derivation of Observation 3

A rise in the costs of treatment increases Δ_h , but decreases Δ_s , the net benefit of treatment to a sick, when treatment costs are identical across states as is assumed. We thus differentiate with respect to Δ_h assuming $\partial \Delta_s / \Delta_h = -1$.

Proof.

$$\frac{\partial VoNI(NT^d)}{\partial \Delta_h} = 0. \quad (\text{A.10})$$

$$\frac{\partial VoNI(T^d)}{\partial \Delta_h} = 0. \quad (\text{A.11})$$

$$\left(\frac{\partial W(NT^d)}{\partial \Delta_h} \right)_{\epsilon^+ = \epsilon^- = 0} = -p_0 - p_0(1 - p_0)\eta(\lambda - 1)\gamma_1. \quad (\text{A.12})$$

$$\left(\frac{\partial W(T^d)}{\partial \Delta_h} \right)_{\epsilon^+ = \epsilon^- = 0} = (1 - p_0) - p_0(1 - p_0)\eta(\lambda - 1)\gamma_1. \quad (\text{A.13})$$

□

Derivation of Observation 4

Proof. To distinguish the effects of an increase in disease severity from changes in the characteristics of treatment we will assume

$$\Delta_s = \alpha r_{NT},$$

$$r_T = \beta r_{NT},$$

$$\Delta_h = (1 - \alpha - \beta)r_{NT},$$

for some $\alpha > 0, \beta > 0$ and $\alpha + \beta < 1$. These assumptions ensure that (a) the benefit-cost ratio of treatment, and (b) the relative effectiveness of treatment remain constant. Using these assumptions it is easy to see that the value of information, be it non-instrumental, instrumental, perfect, or imperfect, is simply a linear function of r_{NT} of the form $\psi \cdot r_{NT}$ where ψ denotes some constant. The change in the value of information will thus be equal to ψ . More importantly, it means that the change in value is positive (negative) if and only if the value itself is positive (negative). □

Derivation of Observation 5

Proof. (i) Speed of testing: variation in τ_1 .

$$\frac{d VoNI(NT^d)}{d \tau_1} = -q^- q^+ (p^+ - p^-) \eta (\lambda - 1) r_{NT} \frac{\partial \gamma_1}{\partial \tau_1} > 0. \quad (\text{A.14})$$

$$\frac{d VoNI(T^d)}{d \tau_1} = -q^- q^+ (p^+ - p^-) \eta (\lambda - 1) r_T \frac{\partial \gamma_1}{\partial \tau_1} > 0. \quad (\text{A.15})$$

$$\begin{aligned} \frac{d VoII(NT^d, i)}{d \tau_1} &= \frac{d VoII(NT^d, n)}{d \tau_1} \\ &= -q^- q^+ \eta (\lambda - 1) [(1 - p^-) \Delta_h + (p^+ - p^-) r_T + p^- \Delta_s] \frac{\partial \gamma_1}{\partial \tau_1} > 0. \end{aligned} \quad (\text{A.16})$$

$$\begin{aligned} \frac{d VoII(T^d, i)}{d \tau_1} &= \frac{d VoII(T^d, n)}{d \tau_1} \\ &= -q^- q^+ \eta (\lambda - 1) [(1 - p^-) \Delta_h + (p^+ - p^-) r_T + p^- \Delta_s] \frac{\partial \gamma_1}{\partial \tau_1} > 0. \end{aligned} \quad (\text{A.17})$$

(ii) Time of testing: variation of τ_1, τ_0 by an equal amount τ , thus $\frac{\partial \tau_1}{\partial \tau} = \frac{\partial \tau_0}{\partial \tau} = 1$.

$$\frac{d VoNI(NT^d)}{d \tau} = -q^- q^+ (p^+ - p^-) \eta (\lambda - 1) r_{NT} \frac{\partial \gamma_1}{\partial \tau_1} > 0. \quad (\text{A.18})$$

$$\frac{d VoNI(T^d)}{d \tau} = -q^- q^+ (p^+ - p^-) \eta (\lambda - 1) r_T \frac{\partial \gamma_1}{\partial \tau_1} > 0. \quad (\text{A.19})$$

$$\begin{aligned} \frac{d VoII(NT^d, i)}{d \tau} &= -q^- q^+ \eta (\lambda - 1) [(1 - p^-) \Delta_h + (p^+ - p^-) r_T + p^- \Delta_s] \frac{\partial \gamma_1}{\partial \tau_1} \\ &\quad + \eta q^+ [\lambda p^+ \Delta_s - (1 - p^+) \Delta_h] \frac{\partial \gamma_0}{\partial \tau_0}. \end{aligned} \quad (\text{A.20})$$

$$\begin{aligned} \frac{d VoII(T^d, i)}{d \tau} &= -q^- q^+ \eta (\lambda - 1) [(1 - p^-) \Delta_h + (p^+ - p^-) r_T + p^- \Delta_s] \frac{\partial \gamma_1}{\partial \tau_1} \\ &\quad + \eta q^- [\lambda (1 - p^-) \Delta_h - p^- \Delta_s] \frac{\partial \gamma_0}{\partial \tau_0}. \end{aligned} \quad (\text{A.21})$$

□

2. APPENDIX TO CHAPTER II

Derivation of m_j as an Optimal Attention Allocation under Cognitive Constraints

Suppose a decision maker (DM) faces the problem of choosing between a finite number of alternatives from the set A . Each alternative is described by a finite vector of attributes $j \in J$. Let x_j^a denote the extent to which alternative a features attribute j . The experienced utility of each alternative $a \in A$ is expressed by

$$u(a) = \sum_j v_j x_j^a, \quad (\text{A.22})$$

where v_j denotes the value the DM ascribes to an additional unit x_j of attribute $j \in J$. Thus the choice problem can be expressed by (A, v) where $A = (x_j^a)_{a \in A, j \in J}$ and $v = (v_j)_{j \in J}$. Suppose a decision-maker (DM) faces cognitive constraints such that he incurs cognitive costs whenever he faces a choice between multi-attribute alternatives. i.e. $|J| \geq 2$. The DM faces no information problem, he perfectly knows the values A and v . He, however, faces problems whenever he needs to integrate this information in order to make a choice. He thus imperfectly considers or takes into account the information, and thus evaluates each alternative by its decision utility given by

$$\tilde{u}(a) = \sum_j [m_j v_j x_j^a + (1 - m_j) v_j x_j^d], \quad (\text{A.23})$$

where $m_j \in [0, 1]$ can be thought of as an attention parameter. $m_j = 1$ denotes full attention, while $m_j = 0$ denotes complete neglect. When neglecting the information in one dimension/attribute, the DM ascribes some value x_j^d to each alternative. Depending on the assumptions one wants to make, this default value may differ. If the DM is “on average right”, one may choose $x_j^d = \bar{x}_j$ where \bar{x}_j is the average value of x_j across the available alternatives. If the DM is pessimistic, $x_j^d = \min_{a \in A} x_j^a$ might be a good assumption. If the DM has some default alternative, $x_j^d = x_j^{default}$ could be reasonable. Regardless of these assumptions, if the DM neglects a dimension, i.e. if $m_j = 0$, he is not able to discriminate between alternatives along this dimension.¹

¹The exact assumptions about x_j^d are irrelevant for the behavior of the DM. For any value of $m_j \in [0, 1]$ and for any modeler’s choice of $(x_j^d)_{j=1, \dots, n}$, a constant $\sum_j (1 - m_j) v_j x_j^d$ is added to the decision utility of each alternative. For a given vector $(m_j)_{j=1, \dots, n}$, this constant is identical across alternatives. It may differ *across* sets of alternatives because m_j is a function of this set, as we are about to derive. Yet, it does not differ across alternatives for a given set of alternatives. It has thus no impact on the desirability of one alternative over another. With no behavioral impact, we drop it in the main section without loss of generality.

Now, let us look at the error from imperfectly considering the dimension j when comparing two alternatives a, b from the set A :

$$\begin{aligned} v_j|x_j^a - x_j^b| - m_j v_j|x_j^a - x_j^b| &= (1 - m_j)v_j(\max_{\{a,b\}} x_j - \min_{\{a,b\}} x_j) \\ &= (1 - m_j)v_j\alpha_j(a, b)\Delta_j, \end{aligned}$$

where $\Delta_j = \max_{a \in A} x_j^a - \min_{a \in A} x_j^a$ is the maximal difference in dimension j between any two alternatives in the set A , and $\alpha_j(a, b) = (\max_{\{a,b\}} x_j - \min_{\{a,b\}} x_j) / \Delta_j$ is the extent to which the actual difference between two alternatives a and b reflects the maximal difference in j in the choice set A . Note that $\alpha_j(a, b) \in [0, 1]$.

An optimal attention allocation will weigh losses from an erroneous representation against losses from incurred cognition costs.² The exact loss function L is, again, a modeling choice. It should include losses from an imperfect problem representation and a loss from exerting cognitive effort. Let us consider the following maximization problem:

$$\max_{m_j \in [0,1]} (-L) = -\frac{1}{2} \sum_j (1 - m_j)E_j - \sum_j \kappa_{r(j)}m_j, \quad (\text{A.24})$$

where E_j represents the loss in dimension j from an imperfect representation of the problem (A, v) . The loss function L is a weighted average of the loss from inattention E_j and of the cost of considering dimension j , $\kappa_{r(j)}$, where the weight on the former is decreasing and the weight of the latter is increasing in m_j . E could have several forms depending on the modeling choice. Consider e.g. the following:

$$E_j = \sum_{a \in A} \sum_{b \in A \setminus a} (1 - m_j)v_j\alpha_j(a, b)\Delta_j\omega_j(a, b) = (1 - m_j)v_j\Delta_j\Omega_j \quad (\text{A.25})$$

with $\sum_{a \in A} \sum_{b \in A \setminus a} \omega_j(a, b) = 1$, and $\Omega_j = \sum_{a \in A} \sum_{b \in A \setminus a} \alpha_j(a, b)\omega_j(a, b)$. Then E_j represents the weighted average of all errors from imperfectly considering dimension j in the choice problem (A, v) . For simplicity, we opt for a weighting $(\omega_j(a, b))_{a \in A, b \in A}$ that puts all weight on the largest errors, i.e. $\omega_j(a, b) > 0 \Leftrightarrow \alpha_j(a, b) = 1$. This implies that $\Omega_j = 1, \forall j \in J$.

The optimal solution of problem (A.24) is then given by

$$m_j = \max \left\{ 0, 1 - \frac{\kappa_{r(j)}}{\mu_j} \right\}, \quad (\text{A.26})$$

²One may argue that the DM's objective is to make the right decision, not to form a correct representation of the world. So, one might want to insert the loss from taking the wrong action into the objective function. Yet, to determine that loss, one needs to know the right action. The problem would amount to choosing the optimal attention with hindsight. Without the information about the correct action, the best thing one can do is to optimize the representation of the world one bases decisions on.

where $\mu_j = v_j \Delta_j$. Hence, whenever $\mu_j < \kappa_{r(j)}$, dimension j is neglected. One could thus interpret μ_j as a measure of importance of dimension j to the DM. If the dimension is important enough compared to the cognitive costs associated with its consideration, it will be taken into account. And, given that it is taken into account ($m_j > 0$), the extent to which a dimension is taken into account rises in the importance of the dimension.

Note that the attention m_j that dimension j receives is crucially determined by the cognitive costs $\kappa_{r(j)}$ associated with its consideration. We want these costs to reflect the rising difficulty of solving increasingly complex problems. It is hence straightforward to assume that considering a single dimension is costless, as there is no complexity involved. Considering a second dimension involves the need to make a first commensurability consideration and is thus associated with some positive cognition costs. Taking into account additional dimensions should become increasingly costly. It thus matters which dimensions are considered “first”. Denote by $r : J \rightarrow \{1, \dots, n\}$ the order in which the attributes are considered. We will henceforth refer to it as the attention hierarchy. Given some place in the hierarchy, each dimension is associated with some cognitive effort cost $\kappa_{r(j)}$.

Given our ideas of rising consideration costs we assume

$$\begin{aligned} \kappa_1 &= 0, \\ \kappa_{r+1} &> \kappa_r, \quad \forall r \in \{1, \dots, n-1\}. \end{aligned} \tag{A.27}$$

The dimension that is considered first, i.e. the dimension that receives rank 1 in the attention hierarchy, is considered without cognitive effort. Taking into account additional dimensions is increasingly costly. The attention hierarchy $r(j)$ is thus crucial for the eventual attention allocation. Consider the following two-step procedure. First, the DM needs to select an attention hierarchy $r : J \rightarrow \{1, \dots, n\}$ that associates each dimension with some consideration costs $\kappa_{r(j)}$. One can think of this as the problem to determine which dimension to consider first, which second, and so on. After assigning a rank to each dimension, the DM solves the above described problem of optimal attention allocation (A.24) given some assignment of consideration costs. The problem can then be solved backwards. Given any assignment $r(j)$, the optimal attention allocation is given by (A.26). Plugging this back into the objective function (A.24) yields

$$\begin{aligned} (-L) &= \sum_{j:m_j>0} \left[-\frac{1}{2} \left(\frac{\kappa_{r(j)}^2}{\mu_j^2} \right) \mu_j - \kappa_{r(j)} + \frac{\kappa_{r(j)}^2}{\mu_j} \right] - \frac{1}{2} \sum_{j:m_j=0} \mu_j \\ &= \sum_{j:m_j>0} \left[\frac{1}{2} \frac{\kappa_{r(j)}^2}{\mu_j} - \kappa_{r(j)} \right] - \frac{1}{2} \sum_{j:m_j=0} \mu_j. \end{aligned}$$

The objective at the first stage, anticipating the result of the second stage, is thus

$$\max_{r(j)} \sum_{j:m_j>0} \left[\frac{1}{2} \frac{\kappa_{r(j)}^2}{\mu_j} - \kappa_{r(j)} \right] - \frac{1}{2} \sum_{j:m_j=0} \mu_j. \quad (\text{A.28})$$

Now we can state the following result:

Proposition A.1. *The optimal assignment $r^*(j)$ satisfies*

$$\mu_j > \mu_{j'} \Rightarrow r(j) < r(j'). \quad (\text{A.29})$$

The optimal assignment assigns higher attention ranks to more important dimensions.

Proof. Under an optimal assignment $r^*(j)$ interchanging the ranks of any two attributes $j, j' \in J$ may not lead to an increase in $(-L)$. Note that the objective function $(-L)$ is additively-separable across dimensions. We may thus confine attention to the parts of the objective function that depend on the two dimensions j and j' .

Suppose $\mu_j = \mu_{j'}$. It is easy to see that interchanging their ranks has no effect on the objective function. We thus only look at cases in which $\mu_j > \mu_{j'}$. Take any ranking r . Under this ranking dimensions j and j' are associated with some cognitive costs $\kappa_{r(j)}, \kappa_{r(j')}$. Denote by $\kappa_h = \max\{\kappa_{r(j)}, \kappa_{r(j')}\}$ and $\kappa_l = \min\{\kappa_{r(j)}, \kappa_{r(j')}\}$. Whichever dimension is assigned κ_l has a higher rank under r . We now show that if r does not assign κ_l to attribute j (the one with strictly higher importance), r cannot be a maximizer for $(-L)$ for some set of cognitive costs $(\kappa_1, \kappa_2, \dots)$ satisfying our assumption (A.27).

Let us distinguish four cases:

- (i) $\kappa_h > \kappa_l > \mu_j > \mu_{j'}$. Both attributes are neglected before and after interchanging the rank. The objective function $(-L)$ is thus invariant to such a change in ranking.
- (ii) $\mu_j > \mu_{j'} > \kappa_h > \kappa_l$. Both attributes are taken into account at the lower rank. However, it is better to assign attribute j the higher rank (and thus κ_l) if

$$\begin{aligned} & \frac{1}{2} \frac{\kappa_l^2}{\mu_j} - \kappa_l + \frac{1}{2} \frac{\kappa_h^2}{\mu_{j'}} - \kappa_h > \frac{1}{2} \frac{\kappa_l^2}{\mu_{j'}} - \kappa_l + \frac{1}{2} \frac{\kappa_h^2}{\mu_j} - \kappa_h \\ \Leftrightarrow & \frac{\kappa_l^2}{\mu_j} + \frac{\kappa_h^2}{\mu_{j'}} > \frac{\kappa_l^2}{\mu_{j'}} + \frac{\kappa_h^2}{\mu_j} \\ \Leftrightarrow & (\kappa_h^2 - \kappa_l^2)\mu_j > (\kappa_h^2 - \kappa_l^2)\mu_{j'} \\ \Leftrightarrow & \mu_j > \mu_{j'}, \end{aligned}$$

which is true by assumption.

- (iii) $\kappa_h > \mu_j > \mu_{j'} > \kappa_l$. Both attributes are considered at the higher rank but neglected at the lower rank. It is better to assign attribute j the higher rank (and thus κ_l) if

$$\begin{aligned} \frac{1}{2} \frac{\kappa_l^2}{\mu_j} - \kappa_l - \frac{1}{2} \mu_{j'} &> \frac{1}{2} \frac{\kappa_l^2}{\mu_{j'}} - \kappa_l - \frac{1}{2} \mu_j \\ \Leftrightarrow \kappa_l^2 \mu_{j'} - \mu_{j'}^2 \mu_j &> \kappa_l^2 \mu_j - \mu_j^2 \mu_{j'} \\ \Leftrightarrow \mu_j \mu_{j'} (\mu_j - \mu_{j'}) &> \kappa_l^2 (\mu_j - \mu_{j'}) \\ \Leftrightarrow \mu_j \mu_{j'} &> \kappa_l^2, \text{ which is true since } \mu_j > \mu_{j'} > \kappa_l. \end{aligned}$$

- (iv) $\mu_j > \kappa_h > \mu_{j'} > \kappa_l$. Attribute j is considered both at the higher and lower rank. Attribute j' is only considered at the higher rank but neglected at the lower rank. Still, it is better to assign attribute j to the higher rank if

$$\frac{1}{2} \frac{\kappa_l^2}{\mu_j} - \kappa_l - \frac{1}{2} \mu_{j'} > \frac{1}{2} \frac{\kappa_l^2}{\mu_{j'}} - \kappa_l + \frac{1}{2} \frac{\kappa_h^2}{\mu_j} - \kappa_h.$$

This indeed holds true, since

$$\begin{aligned} &\frac{1}{2} \kappa_l^2 \left(\frac{1}{\mu_j} - \frac{1}{\mu_{j'}} \right) - \frac{1}{2} \mu_{j'} - \frac{1}{2} \kappa_h^2 \frac{1}{\mu_j} + \kappa_h = \\ &\frac{1}{\mu_j \mu_{j'}} \left[\frac{1}{2} \kappa_l^2 (\mu_{j'} - \mu_j) - \frac{1}{2} \mu_j \mu_{j'}^2 - \frac{1}{2} \kappa_h^2 \mu_{j'} + \kappa_h \mu_j \mu_{j'} \right] = \\ &\frac{1}{2 \mu_j \mu_{j'}} \left[\kappa_l^2 (\mu_{j'} - \mu_j) + (\kappa_h \mu_j \mu_{j'} - \mu_j \mu_{j'}^2) + (\kappa_h \mu_j \mu_{j'} - \kappa_h^2 \mu_{j'}) \right] = \\ &\frac{1}{2 \mu_j \mu_{j'}} \left[\kappa_l^2 (\mu_{j'} - \mu_j) + \mu_j \mu_{j'} (\kappa_h - \mu_{j'}) + \kappa_h \mu_{j'} (\mu_j - \kappa_h) \right] > \\ &\frac{1}{2 \mu_j \mu_{j'}} \left[\kappa_l^2 (\mu_{j'} - \mu_j) + \kappa_h \mu_{j'} (\kappa_h - \mu_{j'}) + \kappa_h \mu_{j'} (\mu_j - \kappa_h) \right] = \\ &\frac{1}{2 \mu_j \mu_{j'}} \left[(\mu_{j'} \kappa_h - \kappa_l^2) (\mu_j - \mu_{j'}) \right] > 0, \end{aligned}$$

in which the first inequality (line 4) holds since $\mu_j > \kappa_h > \kappa_l$, and hence the second term in the bracket is replaced by strictly lower term. The final inequality holds since $\mu_j > \kappa_h > \mu_{j'} > \kappa_l$.

We have now shown that an optimal solution is to set $r(j)$ according to (A.29). One may argue that this rule may be violated for dimensions for which case (i) holds. While this is true, one can counter that (A.29) is optimal for a choice problem (A, v) , and thus for a given vector of $(\mu_j)_{j=1..n}$, for *all* cognitive cost vectors $(\kappa_1, \dots, \kappa_n)$ for which $\kappa_{r+1} > \kappa_r, \forall r = 1, \dots, n$. The optimal assignment (A.29) is thus invariant to changes in the cost vector (e.g. due to changes in cognitive resources for some given choice task). In addition, even if case (i) holds for some dimensions for a given cost vector, it cannot hold for all dimensions as long as $\kappa_1 = 0$. \square

Characteristics of the Attention Function

This section seeks to discuss some characteristics of the attention function, that we derived.

First, more important dimensions receive (weakly) more attention than less important attributes, $\partial m_j / \partial \mu_j \geq 0$. The attention each dimension receives thus depends positively on its value to the decision-maker. In addition to this internal factor, the dimensions' dispersion within the choice set A increases attention. The attention an attribute receives thus depends on the choice environment. It is possible to **attract attention** to a dimension by varying the choice set appropriately.

Second, as one dimension gains importance it may eventually gain rank in the attention hierarchy. As the attention order is strict another dimension must receive a lower rank and thereby loose attention, $\partial m_j / \partial \mu_{j'} \leq 0$, with strict inequality only if $\partial m_{j'} / \partial \mu_{j'} > 0$. It is hence possible to **distract attention** from a dimension. It is noteworthy that this distraction effect only works through the attention hierarchy and is thus discontinuous.

Next, the attention process features **neglect**, or, in Gabaix' terminology, the attention vector is sparse (Gabaix (2011)). Technically, for any decision problem (A, v) there exist vectors of cognitive costs κ satisfying our assumptions such that there exist dimensions $j \in J : m_j = 0$ whenever $|J| \geq 2$. So, for any complex choice problem, i.e. one that involves at least two dimensions, cognitive costs may lead to the neglect of at least one dimension. Similarly, for any vector of cognitive costs κ satisfying our assumptions there exist choice problems $(A, v) : |J| \geq 2$, such that at least one of the dimensions is ignored.

In addition, due to our assumptions on κ , for any choice problem (A, v) there exist attributes $j \in J : m_j > 0$. So, there is **no complete neglect**. As we seek to model the need to simplify a complex choice problem, the DM always considers at least one dimension as this amounts to solving a simple choice problem. This directly implies that complexity costs, as modeled here, will never lead to the choice of an alternative that fares worse than another in all dimensions.

The attention hierarchy is not just implicit. Any two dimensions that are considered receive a **different weight**: $m_j \neq m_{j'}, \forall j, j' \in J : m_j, m_{j'} > 0$. More specifically, for any two dimensions that are considered, the higher ranking dimension receives strictly more attention.

This, together with the impossibility of complete neglect, implies that the attention process *always* features **over- and underweighting**. Let $\bar{m} = \frac{1}{n} \sum_j m_j$. Then for any complex choice problem and cognitive cost vector κ satisfying our assumptions there exist dimensions that are overweighted and dimensions that are underweighted.

Formally, $\forall(A, v, \kappa) : \exists j \in J : m_j > \bar{m}$ and $\exists j \in J : m_j < \bar{m}$. This is important as it implies that under the derived attention function the decision utility of an alternative is not just an affine transformation of experienced utility.

Proof of Proposition II.1

First, let us look at the case in which a monopolist equips his product only with a single quality, say i . In this case, the price is constrained by the fact that the price must rank second in the attention hierarchy: $\mu_i = v_i q_i \geq v_p P = \mu_p \Rightarrow r(i) = 1, r(p) = 2$. Otherwise the decision utility from buying the product would be strictly negative. This yields a maximum price $P^* = v_i x_i / v_p$.³ Then it follows that the profit maximizing choice of the quality level q_i is:

$$q_i^* = \frac{v_i}{v_p c_i}, \quad (\text{A.30})$$

while the resulting profit is

$$\Pi^* = \frac{v_i^2}{2c_i(v_p)^2} = \pi_i. \quad (\text{A.31})$$

The value π_i is the profit that a monopolist can extract from producing a product that features only quality i . It exactly equals the maximum additional profit a monopolist could make by adding this quality to his product under full attention. It is straightforward to see that a monopolist who is confined to produce only a single-quality product will choose the quality with the largest π_i . Since this is always possible the optimal product design always features at least one quality $i : \pi_i = \pi^{(1)}$.

Now, suppose the monopolist contemplates to let the product feature more than a single quality. If the optimal design features more than a single quality, the price dimension must rank first in the attention hierarchy. To see this, note that the price of the optimal single-quality product equals the full value created by the most profitable quality. If introducing more qualities is profitable, then the price of this multi-quality product must exceed the price of the optimal single-quality product. Then the price must also exceed the value of each single quality provided by the multi-quality product. Knowing that $r(p) = 1$ and hence $m_p = 1$, the nonnegativity constraint on the decision value $\tilde{u}(a)$ of the product yields the maximal price, the monopolist can charge for the multi-quality product:

$$P^* = \frac{1}{v_p} \sum_{i \in I} m_i v_i q_i$$

³We assume here that the tie due to $\mu_i = \mu_p$ is broken in favor of the quality dimension. If this is not the case the price is set marginally below $v_i q_i / (v_p)$.

Plugging this back into the profit function allows for maximization with respect to the levels of q_i . It is easy to show that, if $m_i \neq 0$ for some $i \in I$, then $m_i v_i q_i = v_i q_i - \kappa_{r(i)}$ and the optimal level of the quality is $q_i^* = \frac{v_i}{v_p c_i}$. Similarly, it is rather straightforward that if $m_i = 0$, then $q_i^* = 0$ since the production of a quality is costly. Plugging q_i^* back into the profit function yields

$$\Pi^* = \sum_{i \in \mathcal{I}} (\pi_i - \kappa_{r(i)}) \quad (\text{A.32})$$

where \mathcal{I} denotes the subset of qualities of which a positive level is produced under the optimal design. We still need to find the elements of \mathcal{I} , i.e. which qualities are produced under the optimal design of a multi-quality product. Note that for any \mathcal{I} with $|\mathcal{I}| = m \geq 2$ elements, above profit formula gives

$$\Pi^* = \sum_{i \in \mathcal{I}} \pi_i - \sum_{r=2}^m \kappa_r. \quad (\text{A.33})$$

We can therefore conclude that if \mathcal{I} includes m elements, it must be the m most profitable qualities $\{i \mid \exists t \leq m : \pi_i = \pi^{(t)}\}$. How large is m ? It must be true that for $m \geq 2$,

$$\pi^{(m)} - \frac{1}{v_p} \kappa_{m+1} \geq 0$$

and

$$\pi^{(m+1)} - \frac{1}{v_p} \kappa_{m+2} < 0.$$

If that was not the case you could increase profits by either decreasing or increasing \mathcal{I} . This proves part (ii) of Proposition II.1.

Finally, we need to determine whether it is more profitable to offer a single-quality or a multi-quality product. Note that the maximal profit from a multi-quality product with m qualities is given by

$$\sum_{t=1}^m \max \left\{ \pi^{(t)} - \frac{1}{v_p} \kappa_{t+1}, 0 \right\},$$

while the maximal profit from a single-quality product is given by $\pi^{(1)}$. If and only if the latter is larger than the former, i.e. $\sum_{t=2}^m \max \left\{ \pi^{(t)} - \frac{1}{v_p} \kappa_{t+1}, 0 \right\} < \frac{1}{v_p} \kappa_2$, then it is optimal to produce a single-quality product. This proves part (i) of Proposition II.1.

Proof of Proposition II.2

Consider a monopolist whose product features $|\mathcal{I}| \geq 2$ qualities. Further suppose that the optimal level of one of these qualities is below its technological frontier, i.e.

$\exists l \in \mathcal{I} : q_l^* = \frac{v_l}{v_p c_l} < \bar{q}_l$. Let the monopolist offer a second product, called bait good, that has the same level of qualities q_i^b as the initial product, i.e. $q_i^b = q_i, \forall i$. The price of the bait good is set sufficiently high such that it is unattractive to consumers. This introduction of the bait good has no impact on the attention levels with $m_p = 1$ and $m_i = 1 - \kappa_{r(i)}/(v_i q_i^b), \forall i$. However, the profit from the sale of the primary good now has the following form:

$$\Pi^* = P^* - \sum_i c(q_i) = \sum_{i \in \mathcal{I}} \left[\frac{1}{v_p} m_i v_i q_i - \frac{1}{2} c_i q_i^2 \right],$$

with

$$m_i = \begin{cases} \max \{0, 1 - \kappa_{r(i)}/(v_i q_i^b)\} & \text{for } q_i \leq q_i^b, \\ \max \{0, 1 - \kappa_{r(i)}/(v_i q_i)\} & \text{for } q_i > q_i^b. \end{cases}$$

Then it holds that the new optimal level for quality l of the primary good is:

$$q_l^* = \frac{1}{v_p c_l} \left(v_l - \kappa_{r(l)} \frac{1}{q_l^b} \right) < \frac{v_l}{v_p c_l}.$$

Since the original level of the quality was still feasible ($q_l = \frac{v_l}{v_p c_l}$) and the attention allocation under the new design is the same as under the old design, it must hold by revealed preference that the firm makes higher profits with the introduction of the bait good.

Proof of Corollary II.1

We have $\nu_1 \geq \nu_2 > \kappa_3$. That means the more profitable quality also has the higher technological frontier.⁴ Also, no matter which assignment is chosen, both qualities will not be neglected if the bait good levels are set appropriately high.

Note that its always optimal to set $q_i^b = \bar{q}_i$ for the quality i that has the highest rank among all qualities, i.e. rank 2. Accordingly, $m_{12} = 1 - \kappa_2/\nu_1$ and $m_{22} = 1 - \kappa_2/\nu_2$. When receiving the lower attention rank, rank 3, the two qualities receive attention $m_{13} = 1 - \kappa_3/(\min\{\nu_1, \nu_2\}) = 1 - \kappa_3/\nu_2 = m_{23}$. The last equality is straightforward given the insight that it is always optimal to maximize q_i^b if there is no danger to change the attention order. The second equality stems from the fact that q_1^b shall attain the largest value possible at attention rank 3. Given that $q_2^b = \bar{q}_2$ the level of q_1^b is constrained not to exceed ν_2/ν_1 , otherwise quality 1 and 2 would switch rank. Thus

⁴More precisely, the utility that may be delivered through maximizing quality 1 by exploiting its technological boundaries exceeds the utility that may be delivered through maximizing quality 2. For simplicity, we will talk about comparing technological boundaries when comparing ν_1 and ν_2 .

the largest value q_1^b can attain without distracting from quality 2 is ν_2/ν_1 .⁵ Now, since $m_{13} = m_{23}$, $m_{12} \geq m_{22}$ because $\nu_1 \geq \nu_2$, and $\pi_1 > \pi_2$ we conclude that it is strictly more profitable to assign the more profitable quality 1 the higher attention rank.

Proof of Corollary II.2

With $\nu_2 > \nu_1$, we have $m_{12} = 1 - \kappa_2/\nu_1$, $m_{22} = 1 - \kappa_2/\nu_2$, and $m_{13} = m_{23} = 1 - \kappa_3/\min\{\nu_1, \nu_2\} = 1 - \kappa_3/\nu_1$. We want to find a sufficient condition such that

$$\frac{\pi_1}{\pi_2} < \frac{m_{22}^2 - m_{23}^2}{m_{12}^2 - m_{13}^2} \Leftrightarrow \frac{\pi_1 - \pi_2}{\pi_2} < \frac{m_{22}^2 - m_{12}^2}{m_{12}^2 - m_{13}^2}.$$

Consider the right-hand side of the inequality.

$$\frac{m_{22}^2 - m_{12}^2}{m_{12}^2 - m_{13}^2} = \frac{(m_{22} + m_{12})(m_{22} - m_{12})}{(m_{12} + m_{13})(m_{12} - m_{13})} \quad (\text{A.34})$$

$$> \frac{m_{22} - m_{12}}{m_{12} - m_{13}} = \frac{(\nu_2 - \nu_1)\kappa_2}{\nu_2(\kappa_3 - \kappa_2)}. \quad (\text{A.35})$$

The inequality in line 2 holds since $m_{22} > m_{13}$. This gives us the sufficient condition stated in the corollary.

Proof of Corollary II.3

For quality 1 to receive attention rank 2, we must have

$$\frac{\pi_1}{\pi_2} > \frac{m_{22}^2 - m_{23}^2}{m_{12}^2 - m_{13}^2}.$$

With $\nu_2 > \nu_1$, we have $m_{12} = 1 - \kappa_2/\nu_1$, $m_{22} = 1 - \kappa_2/\nu_2$, and $m_{13} = m_{23} = 1 - \kappa_3/\min\{\nu_1, \nu_2\} = 1 - \kappa_3/\nu_1$. Now, if ν_1 converges to ν_2 , m_{12} converges to m_{22} . Therefore

$$\lim_{\nu_1 \rightarrow \nu_2} \frac{m_{22}^2 - m_{23}^2}{m_{12}^2 - m_{13}^2} = 1.$$

Thus, since $\pi_1 > \pi_2$, the required inequality holds for ν_1 sufficiently close to ν_2 .

⁵We make the technical assumption that if $\mu_j = \mu_{j'}$, the firm may select which quality assumes the higher attention rank.

3. APPENDIX TO CHAPTER III

Proof of Proposition III.1

If Assumption III.1 is satisfied, $\mu_p > \mu_{c(i)}$, $\forall i \in F$ and $\mu_p > \mu_{h(i)}$, $\forall i \in \bar{F}$. Since $v_h T_i > v_p c_i \geq v_p(1 - \alpha)c_i$, $\forall \alpha \in [0, 1]$, $\forall i \in \bar{F}$ this also means that $\mu_p > \mu_{c(i)}$, $\forall i \in \bar{F}$. This implies that $r(p) = 1$ and $m_p = 1$. Consider the difference between the difference in experienced utility U and the difference in decision-utility \tilde{U} when $m_p = 1$:

$$U - \tilde{U} = \sum_{i \in F} \pi_i (1 - m_{c(i)}) v_p \alpha c_i + \sum_{i \in \bar{F}} \pi_i (1 - m_{h(i)}) v_h T_i - \sum_{i \in \bar{F}} \pi_i (1 - m_{c(i)}) v_p (1 - \alpha) c_i. \quad (\text{A.36})$$

The first and the second term are strictly positive while the third is strictly negative. If α is (close to) one, thus the co-payment rate is zero (small), the third term is dominated. We want to establish this formally.

Since we want to make a statement about a set of contracts differing in their degree of coverage α , we need to make some assumptions on how the premium P varies with α . Let $P(\alpha)$ denote the premium of a contract offering coverage of α . We assume the premium to be weakly increasing in α : $P(\alpha) \geq P(\alpha')$; $\forall \alpha > \alpha'$. Second, we assume that $P(1) < B$, i.e. the premium for full coverage is affordable.⁶ Denote by $\bar{c} = \max_{i \in I} c_i$ the cost of the most expensive treatment and let $\alpha(\bar{F})$ be the largest $\alpha \in [0, 1]$ such that $B - P(\alpha) = (1 - \alpha)\bar{c}$.⁷ It must be true that $B - P(\alpha) \geq (1 - \alpha)\bar{c}$, $\forall \alpha \in [\alpha(\bar{F}), 1]$, i.e. the decision-maker is able to afford the most expensive treatment, which again implies he is able to afford treatment for all diseases, when buying a contract with coverage $\alpha \geq \alpha(\bar{F})$.

Now, denote by $\underline{\alpha}$ the infimum extent of coverage $\alpha \in [\alpha(\bar{F}), 1]$ such that the following inequality holds for all $i \in \bar{F}$ and all $\alpha \geq \underline{\alpha}$:

$$\begin{aligned} & \pi_i [(1 - m_{h(i)}) v_h T_i - (1 - m_{c(i)}) v_p (1 - \alpha) c_i] \\ & = \min \{ \kappa_{r(h(i))}, \pi_i v_h T_i \} - \min \{ \kappa_{r(c(i))}, \pi_i v_p (1 - \alpha) c_i \} \geq 0. \end{aligned} \quad (\text{A.37})$$

This infimum exists⁸ and is bounded away from 1.⁹ For any level of coverage $\alpha \in (\underline{\alpha}, 1]$,

⁶We continue to abstract from the possibility that a premium is not affordable.

⁷We know - by Tarski's fixed-point theorem - that such an α must exist since $B - P(\alpha)$ is weakly decreasing in α , $(1 - \alpha)\bar{c}$ is continuous and strictly decreasing in α , and $B - P(1) > (1 - 1)\bar{c} = 0$ while $B - P(0) < (1 - 0)\bar{c}$, since $\bar{c} = c_i$ for some $i \in \bar{F}$. We cannot rule out that there is more than one intersection.

⁸The inequalities are satisfied for $\alpha = 1$.

⁹ $\kappa_{r(h(i))}, \pi_i v_h T_i, \kappa_{r(c(i))}$ are all strictly positive since we have established that the premium is the highest-ranking dimension and all lower-ranking dimensions are associated with a strictly positive threshold κ . Hence, there must exist an $\alpha < 1$ such that $\pi_i v_p (1 - \alpha) c_i < \min \{ \kappa_{r(h(i))}, \pi_i v_h T_i, \kappa_{r(c(i))} \}$.

the second and third part sum to a nonnegative number. Since the first part is strictly positive as we established before, we must have $U - \tilde{U} > 0$.

Finally, suppose $\bar{F} \neq \emptyset$ and $U > 0$. It is easy to see that the respective difference in decision utility is negative, $\tilde{U} < 0$, if $m_p = 1$ and the cognitive costs $\kappa_r, r > 1$ are sufficiently large.

Proof of Proposition III.2

We argue that there is scope for profitable undercutting if the difference (III.11) is strictly negative and $m_p = 1$. For as long as $m_p = 1$, the difference in decision utility $\tilde{U}_1 - \tilde{U}_2$ is continuous in P' . Then there exists a third plan with premium P'' such that $P' < P'' < P$ and coverage rate $\alpha'' = \alpha'$ for which $\tilde{U}_3 > \tilde{U}_1$ must hold. Offering this plan attracts customers and is strictly more profitable than the incumbent plan. At the same time it offers strictly lower experienced utility to customers since $P - P'' < \sum_{i \in F \cup A'} (\alpha - \alpha') \pi_i c_i - \sum_{i \in A \setminus A'} \pi_i \alpha c_i$. In the following, we want to show that under the conditions given in Proposition III.2 the difference $\tilde{U}_1 - \tilde{U}_2$ given by (III.11) is strictly negative and $m_p = 1$.

Since insurance is voluntary, the customers have the possibility not to insure. Suppose that, in addition to this outside option, only a single insurance plan that is priced at (or above) the actuarially fair premium is offered and it is demanded by the customers in absence of a second insurance plan. That is the decision utility of buying this first insurance plan (weakly) exceeds the decision utility of remaining uninsured. Construct a second insurance plan by slightly lowering the coverage rate to α' and lowering the premium to $P' = P - \sum_{i \in F \cup A'} (\alpha - \alpha') \pi_i c_i - \sum_{i \in A \setminus A'} \pi_i \alpha c_i$. The difference in decision utility is given by (III.11). Health plan 2 is preferred to plan 1 (and thus also to the outside option) if this difference is negative. Consider the attention parameters m_j of this choice problem. The outside option of “no insurance” is an alternative that is “extreme” on many dimensions. It is the best option in the premium dimension, while the worst one is the first insurance plan as it requires the highest premium payment. Thus, $\mu_p = v_p P$. In the health dimensions for diseases $i \in F$ all available options feature the same consequences. Hence, these dimensions are neglected. In all health outcome dimensions for which the incumbent plan provides access, $h(i), i \in A$, the incumbent plan is the best and no insurance is the worst option: $\mu_{h(i)} = \pi_i v_h T_i$. In the co-payment dimensions for diseases with affordable treatment, $c(i), i \in F$ the best option is plan 1 (lowest co-payment) while the worst option is no insurance (full payment): $\mu_{c(i)} = v_p \pi_i (1 - (1 - \alpha)) c_i$. In the co-payment dimensions for diseases $i \in A'$, the best option is no insurance (no expenditure) and the worst is plan 2 (highest co-payment): $\mu_{c(i)} = v_p \pi_i (1 - \alpha' - 0) c_i$. Finally, in the co-payment dimensions for diseases

$i \in A \setminus A'$, the best option is no insurance (no expenditure) and the worst is plan 1 (highest co-payment): $\mu_{c(i)} = v_p \pi_i (1 - \alpha - 0) c_i$. The attention parameters are thus given by

$$\begin{aligned} m_p &= \max \{0, 1 - \kappa_{r(p)} / (v_p (P - 0))\}, \\ m_{h(i)} &= 0, \forall h(i) : i \in F, \\ m_{h(i)} &= \max \{0, 1 - \kappa_{r(h(i))} / (\pi_i v_h T_i)\}, \forall h(i) : i \in A, \\ m_{c(i)} &= \max \{0, 1 - \kappa_{r(c(i))} / (v_p \pi_i (1 - (1 - \alpha) c_i))\}, \forall c(i), i \in F, \\ m_{c(i)} &= \max \{0, 1 - \kappa_{r(c(i))} / (v_p \pi_i (1 - \alpha' - 0) c_i)\}, \forall c(i), i \in A', \\ m_{c(i)} &= \max \{0, 1 - \kappa_{r(c(i))} / (v_p \pi_i (1 - \alpha - 0) c_i)\}, \forall c(i), i \in A \setminus A'. \end{aligned}$$

Now suppose that α and $\alpha' < \alpha$ are such that $A = A'$. Then the second term in (III.11) vanishes. A sufficient condition for (III.11) to be negative is then $m_p = 1$. If assumption III.2 is satisfied, then $\mu_p > \mu_j, \forall j \neq p$. This again implies that $1 = m_p > m_{c(i)}, \forall i \in F \cup A'$. As a result, the decision utility of the low-quality plan 2 will exceed the decision utility of the high-quality plan 1, although the high-quality plan offers higher experienced utility.

Now consider the case when any reduction in coverage entails a loss in access $A' \subset A, \forall \alpha' < \alpha$. That is, we assume the incumbent policy offers some coverage α such that $A(\alpha') \subset A(\alpha), \forall \alpha' < \alpha$. In this case, profitable undercutting cannot always work out. Since we assume that the first health plan is demanded in the absence of health plan 2, we know that $m_{h(i)} > 0$ for some $i \in A$. Thus, even if we maintain the assumption that $v_p P > \max_{i \in A} \pi_i v_h T_i$ such that the premium dimension receives full attention, we know that the access value of at least some diseases must be sufficiently large that they make insurance desirable, even when underappreciated. Thus there must exist some $i \in A$ for which the removal of access is noticed and sufficiently undesirable to make undercutting infeasible. However, it is not guaranteed that these are exactly the ones that A' lacks. There can be diseases $i \in A$ for which undercutting an incumbent policy with coverage α is feasible, in particular if the number of diseases covered $|F \cup A|$ is large. Note that since an outside option is available the decision-maker may consider health dimensions $h(i) : i \in A'$ in which the two insurance plans do not differ while paying less or no attention to dimensions in which there are differences between the plans $h(i) : i \in A \setminus A'$. If he happens to neglect exactly the health dimension $h(i) : i \in A \setminus A'$ then the loss in access due to the slight reduction of coverage remains unrecognized. Thus, if $m_{h(i)} = 0$ for $i \in A \setminus A'$ and if assumption III.2 is satisfied for plan 1 and thus $m_p = 1$, then $\tilde{U}_1 - \tilde{U}_2$ is negative and, hence, profitable undercutting is feasible.

Proof of Proposition III.3

Denote by ψ a dimension $j \in \{p, c(i) : i \in F, h(i) : i \in \bar{F}\}$ for which $\mu_j > 0$, yet $m_j = 0$. If $\psi \neq p$, denote by ι the disease $i \in I$ of which ψ is either a monetary or health consequence. Denote by $\bar{g} \in \operatorname{argmax}_{g \in \Gamma} u(g, \psi)$ one of the available alternatives with maximal utility in dimension ψ . Denote by $\underline{g} \in \operatorname{argmin}_{g \in \Gamma} u(g, \psi)$ one of the alternatives with minimal utility in dimension ι . Construct a plan g' such that $u(g', \psi) = u(\underline{g}, \psi)$ and $u(g', j) = u(\bar{g}, j)$, $\forall j \in B \setminus \psi$. More precisely, if $\psi = p$, set the price of g' equal to the price of the most expensive plan available, \underline{g} , and set the levels of coverage α'_i equal to the levels of coverage α_i of the cheapest plan available, \bar{g} . Alternatively, if $\psi \neq p$ construct g' by equating the level of coverage for disease ι to the lowest level of coverage for ι available (under \underline{g}) while equating the price P' , and the levels of coverage $\alpha'_i, i \in I \setminus \iota$ to the levels provided by the plan \bar{g} that offers highest coverage of ι . It is easy to see that extending Γ by g' does not change the attention allocation since the range of utility μ_j in each dimension remains unchanged. It follows that dimension ψ remains neglected if g' is included in the choice set. g' is constructed to be equal to \bar{g} in all dimensions but ψ , in which it is inferior, hence \bar{g} dominates g' . But since ψ is neglected by the decision-maker, he is indifferent between these two alternatives.

Proof of Corollary III.1

The stated condition requires that there exists a dimension $j \in \{p, c(i) : i \in F, h(i) : i \in \bar{F}\}$, call it ψ , in which an alternative g^* that would be chosen from Γ holds an advantage over some other available alternative, yet this advantage is neglected. In this case the dominated alternative g' is constructed as in the proof of Proposition III.3 by replacing \bar{g} with g^* . The newly constructed alternative g' will be dominated by g^* , yet the decision-maker will be indifferent between g^* and g' . Since g^* is maximizing decision-utility among all alternatives from the choice set Γ , and, since the attention allocation remains unchanged, also from the choice set $\Gamma \cup g'$, it follows that g' must also maximize decision-utility among all alternatives from $\Gamma \cup g'$. It follows that the decision-maker would be willing to choose g' despite it being dominated.

Proof of Proposition III.4

We consider insurance plans that fully pay the treatment cost for the diseases they cover. The difference between the decision utility of buying insurance covering the nonempty set of diseases S and the decision utility of not buying insurance is then given by $\tilde{U}(S) = \sum_{i \in S} m_{i,S} b_i - m_{p,S} v_p P$. $m_{j,S}$ denotes the attention a dimension j receives when the choice set is given by $\Gamma = \{(S, P), (\emptyset, 0)\}$. b_i denotes the benefit

of having insurance covering the full treatment cost for disease i . That means, $b_i = \pi_i c_p c_i, \forall i \in F$ and $b_i = \pi_i v_H T_i, \forall i \in \bar{F}$.¹⁰

Let $W(S) = \max \left\{ P : \tilde{U}(S) \geq 0 \right\}$ be the maximum willingness-to-pay for an insurance that fully covers treatment costs of diseases $i \in S$. We make the following technical assumption. If $\mu_p = \mu_j, j \neq p$, then $r(p) > r(j)$. That is, if the premium dimension ties with another dimension, this other dimension gains higher rank in the attention hierarchy. This assumption ensures that the maximum premium $P : \tilde{U} \geq 0$ always exists.

We now establish that $W(S) = \max \left\{ \max_{i \in S} b_i / v_p, \sum_{i \in S} \bar{m}_{i,S} b_i / v_p \right\}$ where $\bar{m}_{j,S}$ is the attention parameter of dimension j if the attention rank of the premium is bound to be $r(p) = 1$, while the remaining ranks are determined as usual according to $\mu_j > \mu_{j'} \Rightarrow r(j) < r(j')$.

It is easy to see that $W(S) \geq \max_{i \in S} b_i / v_p$. Suppose not and consider $P < \max_{i \in S} b_i / v_p$ and let j be disease $i \in S$ with maximum expected benefit b_i . Then $\mu_p \leq \mu_j$ and thus $m_{p,S} < m_{j,S}$. This suffices to let $\tilde{U}(S) > 0$. As this holds true for all levels of $P \leq \max_{i \in S} b_i / v_p$, the premium could be increased up to the amount $\max_{i \in S} b_i / v_p$ with \tilde{U} remaining strictly positive. Now, if $W(S) > \max_{i \in S} b_i / v_p$ then the premium must rank first in the attention hierarchy as $\mu_p = v_p P > \max_{i \in S} b_i = \max_{i \in S} \mu_i$ and thereby $m_{S,p} = 1$. Then, from $\tilde{U} = 0$ one can easily verify that $W(S) = \frac{1}{v_p} \sum_{i \in S} \bar{m}_{i,S} b_i$ must be true since $m_{j,S} = \bar{m}_{j,S}$.

Next, we show that $W(C) < W(A) + W(B)$ for any disjoint, nonempty sets of diseases A, B and $C = A \cup B$.

First, suppose that $W(C) = \max_{i \in C} b_i / v_p$. Then $W(C) < \max_{i \in A} b_i / v_p + \max_{i \in B} b_i / v_p \leq W(A) + W(B)$.

Second, suppose that $W(C) = \frac{1}{v_p} \sum_{i \in C} \bar{m}_{i,C} b_i$. Then

$$W(C) < \frac{1}{v_p} \sum_{i \in A} \bar{m}_{i,A} b_i + \sum_{i \in B} \bar{m}_{i,B} b_i \leq W(A) + W(B). \quad (\text{A.38})$$

The second inequality holds by definition of $W(\cdot)$. The first strict inequality is due to the fact that adding further benefit dimensions to the choice problem can never increase the attention rank of (and thus the attention attributed towards) the previous benefit dimensions. Moreover, when “merging” two insurance plans into one comprehensive plan, some of the benefit dimensions must lose rank as the attention hierarchy is strict. In contrast, as the willingness-to-pay for the comprehensive insurance will be at least as high as the willingness-to-pay for each of the individual insurances the premium

¹⁰As the two options only differ in either the health consequence or the monetary consequence in case of a disease, we refrain from differentiating between subscripts $c(i)$ and $h(i)$ for the attention parameters m_j .

dimension cannot lose rank through the merger.

We now show that at least one benefit dimension receives strictly less attention which implies the first strict inequality in (A.38). First, suppose $W(A) = \frac{1}{v_p} \sum_{i \in A} \bar{m}_{i,A} b_i$ and $W(B) = \frac{1}{v_p} \sum_{i \in B} \bar{m}_{i,B} b_i$, i.e. that the premium dimension ranks first for both insurance plans before the merger. Consider for each of the two plans that are merged the *benefit* dimension i that ranks highest in the attention hierarchy. For both of these dimensions, call them a and b , it must be that $\bar{m}_{a,A} > 0$ and $\bar{m}_{b,B} > 0$. Otherwise, e.g. if $\bar{m}_{a,A} = 0$, then $\frac{1}{v_p} \sum_{i \in A} \bar{m}_{i,A} b_i = 0 \neq W(A)$. One of the dimensions must lose rank through the merger since it cannot be that both maintain the rank two as the attention hierarchy is strict. From the definition of the attention parameters m_j it is easy to see that: If (and only if) a dimension receives attention, i.e. $m_j > 0$, then a loss in rank implies a loss in attention (m_j). Therefore, as both highest-ranking benefit dimensions were considered before the merger and one of them loses rank, say a , it must be that this dimension receives strictly less attention, so that $m_{a,C} < m_{a,A}$. As all benefit dimensions receive weakly less attention and there is at least one dimension that receives strictly less attention, it must be that $W(C) < \frac{1}{v_p} \sum_{i \in A} \bar{m}_{i,A} b_i + \sum_{i \in B} \bar{m}_{i,B} b_i$. Consider, on the other hand, the cases in which $W(A) = \max_{i \in A} b_i/v_p$, or $W(B) = \max_{i \in B} b_i/v_p$, or both. W.l.o.g. suppose $W(A) = \max_{i \in A} b_i/v_p = b_a/v_p$, where we again call a the b_i -maximal disease in set A . Since we consider the case in which $W(C) = \frac{1}{v_p} \sum_{i \in C} m_{i,C} b_i$, we know that dimension a ranks first in the attention hierarchy before the merger, while the premium ranks first after the merger. Thus, dimension a must have lost rank through the merger, and since $m_{a,A} > 0$, we know that this loss in rank was accompanied by a loss in attention m_a . Again, since all benefit dimensions receive weakly less attention and there is at least one that receives strictly less attention, we can conclude that $W(C) < \frac{1}{v_p} \sum_{i \in A} \bar{m}_{i,A} b_i + \sum_{i \in B} \bar{m}_{i,B} b_i$. As we have shown that $W(C) < W(A) + W(B)$ for arbitrary non-empty, disjoint sets A, B and $C = A \cup B$, the proposition follows.

Proof of Proposition III.5

In the proposed equilibrium candidate, each firm earns zero profit. Also, the maximum premium charged by any firm in the market is $\sum_{i \in I} \pi_i c_i$, the premium charged for the comprehensive insurance $S = I$, while the minimum premium payment is realized in the outside option of no insurance. If $\sum_{i \in I} \pi_i c_i > b_i \forall i \in \bar{F}$ then $\mu_p > \mu_i \forall i \in \bar{F}$. Since $\sum_{i \in I} \pi_i c_i > b_i = \pi_i c_i \forall i \in F$ always holds, we conclude that $r(p) = 1$ and $m_p = 1$. Hence, customers always recognize premium differences. In this case, no firm can profitably deviate. To show this, we first want to establish that no deviation can

change the attention allocation $m_j, j \in I \cup p$ prevailing in the equilibrium candidate. For the benefit dimensions $j \in I$, we have $\mu_j = b_j - 0 = b_j$ since there are firms that offer plans including insurance against j , thus offering b_j , and there are options, specifically the outside option, that do not include this benefit. No deviation by a single firm could change μ_j since there is neither a way to offer lower utility in dimension j than zero nor a way to offer higher utility than b_j in dimension j . Now, consider the premium dimension p . The maximum utility in this dimension is zero given by the outside option. The lowest utility is set by the comprehensive plan that covers all diseases. It is thus possible to further increase μ_p if the deviant plan would include a premium payment larger than $\sum_I \pi_i c_i$. However, since the premium dimension already ranks first in the attention hierarchy in the equilibrium candidate, $r(p) = 1$, and thereby $m_p = 1$, such a deviation could not further increase the attention allocated to the premium dimension. We conclude that the attention allocation induced in the equilibrium candidate cannot be changed by any deviant plan. We now want to argue that no deviant plan could earn a strictly positive profit. Suppose the contrary, i.e. there exists a plan (\bar{S}, \bar{P}) $\bar{S} \subseteq I$ that is not in the set of plans offered in our equilibrium candidate to which one of the firms could deviate and earn a strictly positive profit. Then $\bar{P} > \sum_{i \in \bar{S}} \pi_i c_i$ must hold for this plan. However, in the equilibrium there exists a firm offering the plan $(\bar{S}, \sum_{i \in \bar{S}} \pi_i c_i)$, i.e. a firm offering the same insurance benefits at an actuarially fair premium. Since premium differences are recognized because $m_p = 1$, no customer would choose the deviant plan. The deviant firm would thus make zero profits, a contradiction.

Proof of Proposition III.6

In the equilibrium of Proposition III.5, the attention hierarchy is given by

$$\begin{aligned} r(p) &= 1, \\ r(i) &= |\{i' \in I : b_{i'} > b_i\}| + 1, \end{aligned}$$

since $\mu_i = b_i \forall i \in I$. This means that the diseases with the highest insurance benefits receive more attention. This implies $m_i = \max\left\{0, 1 - \frac{\kappa_{r(i)}}{b_i}\right\}$, $\forall i \in I$. Customers will purchase the plan with the highest decision utility given this attention allocation.¹¹

We want to establish that the plan maximizing decision utility cannot include any insurance of diseases $i \in I$ with $b_i - v_p \pi_i c_i < \kappa_{r(i)}$. Suppose otherwise and $\exists i \in S^* : b_i - v_p \pi_i c_i < \kappa_{r(i)}$. Then consider the plan S', P' with $S' = S^* \setminus i$ and $P' = \sum_{i \in S'} \pi_i c_i$. The difference in decision-utility between the two plans is given by

$$\tilde{U}(S^*) - \tilde{U}(S') = m_i b_i - m_p v_p \pi_i c_i = \max\{b_i - \kappa_{r(i)}, 0\} - v_p \pi_i c_i < 0.$$

¹¹We will consider the outside option of no insurance as the “plan” $(\emptyset, 0)$.

The plan (S', P') would thus give higher decision utility to the consumer, a contradiction to (S^*, P^*) maximizing decision utility.

Now, suppose in contrast to the claim that $\exists i \in I \setminus S^* : b_i - v_p \pi_i c_i \geq \kappa_{r(i)}$. Construct the plan (S', P') with $S' = S^* \cup i$. The difference in decision utility between the two plans is given by

$$\tilde{U}(S^*) - \tilde{U}(S') = m_p v_p \pi_i c_i - m_i b_i = v_p \pi_i c_i - (b_i - \kappa_{r(i)}) \leq 0,$$

where the second equality holds since $b_i - v_p \pi_i c_i \geq \kappa_{r(i)} \Rightarrow b_i > \kappa_{r(i)}$. Hence, customers would choose (S', P') over (S^*, P^*) . Again, we have a contradiction.

Proof of Claim: If only plans covering $i \in \bar{F}$ are offered in the market, profitable deviations may occur.

Suppose there are only plans available that cover diseases $i \in \bar{F}$, the most comprehensive being $(\bar{F}, \sum_{\bar{F}} \pi_i c_i)$. Further suppose that $\sum_{\bar{F}} \pi_i c_i > b_i, \forall i \in \bar{F}$, such that the premium dimension ranks first again: $r(p) = 1$. The attention hierarchy amongst dimensions on which the available options differ will be given by

$$r(p) = 1, \tag{A.39}$$

$$r(i) = |\{i' \in \bar{F} : b_{i'} > b_i\}| + 1 \tag{A.40}$$

Let (S^*, P^*) denote the plan customers would choose from the available set. Denote by $\iota = \min_{i \in S^*} b_i$ the disease with the lowest insurance benefit covered under this plan. Assume that $b_\iota < \kappa_{r(\iota)+1}$, i.e. this benefit would be neglected at the next lower attention rank. Now, suppose there exists a disease $i \in \bar{F}$ such that $\sum_{\bar{F}} \pi_i c_i > b_i > b_j, \forall j \in \bar{F}$. Construct an additional plan that includes coverage of i yet excludes coverage of ι and suppose it is priced actuarially fair. The difference in decision utility between this new plan and the plan (S^*, P^*) is given by

$$m_i^D b_i - m_\iota^D b_\iota - m_p^D v_p (\pi_i c_i - \pi_\iota c_\iota),$$

where $m_j^D, j = p, i, \iota$ denotes the attention paid to dimension j after the inclusion of the newly constructed plan into the choice set. Note first that since $b_i > b_j, \forall i \in \bar{F}$ it must be true that $m_i^D = 1 - (\kappa_2/b_i)$. Second, $m_\iota^D = 0$. After the inclusion of the new plan there is now an additional dimension, i , that incorporates a larger utility difference. Hence, dimension ι loses a rank in the hierarchy and is now neglected. Finally, since $m_p = 1$ before, we now must have $m_p^D \leq 1$. Thus,

$$m_i^D b_i - m_\iota^D b_\iota - m_p^D v_p (\pi_i c_i - \pi_\iota c_\iota) \geq b_i - \kappa_2 - v_p (\pi_i c_i - \pi_\iota c_\iota) \tag{A.41}$$

$$= v_p \pi_i c_i - \kappa_2 - v_p (\pi_i c_i - \pi_\iota c_\iota) \tag{A.42}$$

$$= v_p \pi_\iota c_\iota - \kappa_2. \tag{A.43}$$

Therefore, if $v_p \pi_i c_i > \kappa_2$, the customers strictly prefer the new plan to (S^*, P^*) . This would allow a small increase in the premium of the new plan without breaking the strict preference. Hence, the newly constructed plan could attract the whole market while making a positive profit.

Bibliography

- Abaluck, J. and J. Gruber (2011). Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. *American Economic Review*, 1180–1210.
- Abeler, J., A. Falk, L. Goette, and D. Huffman (2011). Reference points and effort provision. *The American Economic Review* 101, 470–492(23).
- Barigozzi, F. and R. Levaggi (2008). Emotions in physician agency. *Health Policy* 88, 1–14.
- Barigozzi, F. and R. Levaggi (2010). Emotional decision-makers and anomalous attitudes toward information. *Journal of Risk and Uncertainty* 40, 255–280.
- Bell, D. E. (1985). Disappointment in decision-making under uncertainty. *Operations Research* 33(1), 1–27.
- Benartzi, S. and R. Thaler (2002). How Much is Investor Autonomy Worth? *The Journal of Finance* 57(4), 1593–1616.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics* 127(3), 1243–1285.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and Consumer Choice. *Journal of Political Economy* 121(5), 803–843.
- Caplin, A. and K. Eliaz (2003). Aids policy and psychology: a mechanism-design approach. *The RAND Journal of Economics* 34(4), 631–646.
- Caplin, A. and J. Leahy (2004). The supply of information by a concerned expert. *The Economic Journal* 114, 487–505.
- Chetty, R., A. Looney, and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *The American Economic Review* 99(4), 1145–1177.
- Crawford, V. P. and J. Meng (2011). New york city cabdrivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *The American Economic Review* 101(5).

- Cubanski, J. and P. Neuman (2007). Status Report on Medicare Part D Enrollment in 2006: Analysis of Plan-Specific Market Share and Coverage. *Health Affairs* 26, w1–w12.
- Deblonde, J., P. D. Koker, F. F. Hamers, J. Fontaine, S. Luchters, and M. Temmerman (2010). Barriers to hiv testing in europe: a systematic review. *The European Journal of Public Health*, 1–11.
- Drolet, A., I. Simonson, and A. Tversky (2000). Indifference Curves that Travel with the Choice Set. *Marketing Letters* 11(3), 199–209.
- Eliasz, K. and A. Schotter (2007). Experimental testing of intrinsic preferences for noninstrumental information. *The American Economic Review* 97(2), 166–169.
- Eliasz, K. and A. Schotter (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior* 70(2), 304–324.
- Eliasz, K. and R. Spiegler (2006). Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior* 56(1), 87–104.
- Eliasz, K. and R. Spiegler (2011). On the strategic use of attention grabbers. *Theoretical Economics* 6(1), 127–155.
- Fang, H., M. Keane, and D. Silverman (2006). Sources of Advantageous Selection: Evidence from the Medigap Insurance Market. *NBER Working Paper 12289*, 1–44.
- Feldman, R. and B. Dowd (1991). A New Estimate of the Welfare Loss of Excess Health Insurance. *The American Economic Review* 81(1), 297–301.
- Feldstein, M. (1973). The Welfare Loss of Excess Health Insurance. *Journal of Political Economy* 81(2), 251–280.
- Frank, R. and K. Lamiraud (2008). Choice, Price Competition and Complexity in Markets for Health Insurance. *NBER Working Paper 13817*, 1–44.
- Gabaix, X. (2011). A Sparsity-Based Model of Bounded Rationality. NBER Working Paper No. 16911.
- Gabaix, X. and D. Laibson (2006). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *The Quarterly Journal of Economics* 121(1), 505–540.
- Geer, K. P., M. E. Ropka, W. F. Cohn, S. M. Jones, and S. Miesfeldt (2001). Factors influencing patients' decisions to decline cancer genetic counseling services. *Journal of Genetic Counseling* 10(1), 25–40.

- Gill, D. and V. Prowse (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review* 102(1), 469.
- Greenwald, J. L., G. R. Burstein, J. Pincus, and B. Branson (2006). A rapid review of rapid hiv antibody tests. *Current Infectious Disease Reports* 8, 125–131.
- Handel, B. (2013). Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *The American Economic Review* 103(7), 2643–2682.
- Heidhues, P. and B. Köszegi (2008). Competition and price variation when consumers are loss averse. *The American Economic Review* 98(4), 1245–1268.
- Heidhues, P. and B. Köszegi (2010). Regular prices and sales. *Working Paper*.
- Heiss, F., A. Leive, D. McFadden, and J. Winter (2012). Plan Selection in Medicare Part D: Evidence from Administrative Data. *NBER Working Paper 18166*, 1–47.
- Heiss, F., D. McFadden, and J. Winter (2006). Who Failed to Enroll in Medicare Part D, and Why? Early Results. *Health Affairs* 25, w344–w354.
- Heiss, F., D. McFadden, and J. Winter (2010). Mind the Gap! Consumer Perceptions and Choices of Medicare Part D Prescription Drug Plans. In D. Wise (Ed.), *Research Findings in the Economics of Aging*, pp. 413–481. The University of Chicago Press.
- Herne, K. (1997). Decoy alternatives in policy choices: Asymmetric domination and compromise effects. *European Journal of Political Economy* 13(3), 575–589.
- Herweg, F., D. Müller, and P. Weinschenk (2010). Binary payment schemes: Moral hazard and loss aversion. *The American Economic Review* 100(5), 2451–2477(27).
- Huber, J., J. Payne, and C. Puto (1982). Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis. *The Journal of Consumer Research* 9(1), 90–98.
- Iyengar, S. and E. Kamenica (2010). Choice Proliferation, Simplicity Seeking, and Asset Allocation. *Journal of Public Economics* 94, 530–539.
- Johnson, E., J. Hershey, J. Meszaros, and H. Kunreuther (1993). Framing, Probability Distortions, and Insurance Decisions. *Journal of Risk and Uncertainty* (7), 35–51.
- Johnson, J. and D. Myatt (2006). On the Simple Economics of Advertising, Marketing, and Product Design. *The American Economic Review* 96(3), 756–784.
- Kahneman, D. (2000). Experienced Utility and Objective Happiness: A Moment-Based Approach. In D. Kahneman and A. Tversky (Eds.), *Choices, Values, and*

- Frames*, pp. 673–692. Cambridge University Press.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Kahneman, D. and A. Tversky (1984). Choices, Values, and Frames. *American Psychologist* 39(4), 341–350.
- Kamenica, E. (2008). Contextual Inference in Markets: On the Informational Content of Product Lines. *The American Economic Review* 98(5), 2127–2149.
- Karlsson, N., G. Loewenstein, and D. Seppi (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty* 38, 95–115.
- Kassler, W. J., B. A. Dillon, C. Haley, W. K. Jones, and A. Goldman (1997). On-site, rapid hiv testing with same-day results and counseling. *AIDS* 11, 1045–1051.
- Kőszegi, B. (2003). Health anxiety and patient behavior. *Journal of Health Economics* 22, 1073–1084.
- Kőszegi, B. (2006). Emotional agency. *Quarterly Journal of Economics* 99(3), 909–936.
- Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4), 1133–1165.
- Kőszegi, B. and M. Rabin (2007). Reference-dependent risk attitudes. *The American Economic Review* 97(4), 1047–1073.
- Kőszegi, B. and M. Rabin (2009). Reference-dependent consumption plans. *The American Economic Review* 121(1), 121–155.
- Kőszegi, B. and A. Szeidl (2011). A Model of Focusing in Economic Choice. Working Paper.
- Kivetz, R., O. Netzer, and V. Srinivasan (2004). Alternative Models for Capturing the Compromise Effect. *Journal of Marketing Research* 91, 237–257.
- Kling, J. R., S. Mullainathan, E. Shafrir, L. C. Vermeulen, and M. V. Wrobel (2012). Comparison Friction: Experimental Evidence from Medicare Drug Plans. *The Quarterly Journal of Economics* 127, 199–235.
- Kőszegi, B. and A. Szeidl (2013). A Model of Focusing in Economic Choice. *The Quarterly Journal of Economics* 128(1), 53–104.
- Lerman, C., J. Seay, A. Balshem, and J. Audrain (1995). Interest in genetic testing among first-degree relatives of breast cancer patients. *American Journal of Medical Genetics* 57, 385–392.

- Liebman, J. and R. Zeckhauser (2008). Simple Humans, Complex Insurance, Subtle Subsidies. *NBER Working Paper 14330*, 1–31.
- Loewenstein, G. F., E. U. Weber, C. K. Hsee, and N. Welch (2007). Risk as feelings. In G. Loewenstein (Ed.), *Exotic Preferences: Behavioral Economics and Human Motivation*, pp. 565–612. Oxford University Press.
- Loomes, G. and R. Sugden (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies* 53(2), 271–282.
- Matthey, A. (2008). Yesterday’s expectation of tomorrow determines what you do today: The role of reference-dependent utility from expectations. *Jena Economic Research Paper No. 2008-003*.
- McFadden, D. (1999). Rationality for Economists? *Journal of Risk and Uncertainty* 19(1), 73–105.
- Neuman, P. and J. Cubanski (2009). Medicare Plan D Update - Lessons Learned and Unfinished Business. *The New England Journal of Medicine* 361, 406–414.
- Neumann, P. J., J. K. Hammitt, C. Mueller, H. M. Fillit, J. Hill, N. A. Tetteh, and K. S. Kosik (2001). Public attitudes about genetic testing for alzheimer’s disease. *Health Affairs* 20(5), 252–264.
- Norman, D. (2007). Simplicity is Highly Overrated. *interactions* 14(2), 40–41.
- Nyman, J. (2003). *The Theory of Demand for Health Insurance*. Stanford University Press.
- Osborne, M. and A. Rubinstein (1998). Games with Procedurally Rational Players. *The American Economic Review* 88(4), 834–847.
- Panidi, K. (2008). Why do we avoid doctors? the view from behavioral economics standpoint. Available at SSRN: <http://ssrn.com/abstract=1150328>.
- Post, T., M. J. van den Assem, G. Baltussen, and R. H. Thaler (2008). Deal or no deal? decision making under risk in a large-payoff game show. *The American Economic Review* 98(1), 38–71.
- Radner, R. and J. Stiglitz (1984). A nonconcavity in the value of information. In M. Boyer and R. Kihlstrom (Eds.), *Bayesian Models in Economic Theory*, pp. 33–52. Elsevier Science Publishers B.V.
- Rubinstein, A. (1993). On Price Recognition and Computational Complexity in a Monopolistic Model. *Journal of Political Economy* 101(3), 473–484.
- Rust, R., D. Thompson, and R. Hamilton (2006). Defeating Feature Fatigue. *Harvard Business Review* 48(2), 37–47.

- Schram, A. and J. Sonnemans (2011). How Individuals Choose Health Insurance: An Experimental Analysis. *European Economic Review* 55, 799–819.
- Simonson, I. (1989). Choice Based on Reasons: The Case of Attraction and Compromise Effects. *The Journal of Consumer Research* 16(2), 158–174.
- Simonson, I. and A. Tversky (1992). Choice in Context: Tradeoff Contrast and Extremeness Aversion. *Journal of Marketing Research* 29(3), 281–295.
- Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50, 665–690.
- Sinaiko, A. and R. Hirth (2011). Consumers, Health Insurance, and Dominated Choices. *Journal of Health Economics* 30, 450–457.
- Spiegler, R. (2006a). Competition over agents with boundedly rational expectations. *Theoretical Economics* 1(2), 207–231.
- Spiegler, R. (2006b). The Market for Quacks. *The Review of Economic Studies* 73(4), 1113–1131.
- Sydnor, J. (2010). (Over)insuring Modest Risks. *American Economic Journal: Applied Economics* 2(4), 177–199.
- Thompson, D., R. Hamilton, and R. Rust (2005). Feature Fatigue: When Product Capabilities Become Too Much of a Good Thing. *Journal of Marketing Research* 42, 431–442.
- Vikander, N. (2010). Targeted Advertising and Social Status. Working Paper.
- Weiser, S. D., M. Heisler, K. Leiter, F. P. de Korte, S. Tlou, S. DeMonner, N. Phaladze, D. R. Bangsberg, and V. Iacopino (2001). Routine hiv testing in botswana: A population-based study on attitudes, practices, and human rights concerns. *PLoS Med* 7(3), e261.
- Wernerfelt, B. (1995). A Rational Reconstruction of the Compromise Effect: Using Market Data to Infer Utilities. *Journal of Consumer Research* 21(4), 627–633.
- Zapka, J. G., A. Stoddard, M. Zorn, J. McCusker, and K. H. Mayer (1991). Hiv antibody test result knowledge, risk perceptions and behavior among homosexually active men. *Patient Education and Counseling* 18, 9–17.
- Zhou, J. (2007). Advertising, Misperceived Preferences, and Product Design. Working Paper.