Optimal Revelation of Life-Changing Information*

Nikolaus Schweizer and Nora Szech

August 2017

Abstract

Information about the future may be instrumentally useful, yet scary. For example, many patients shy away from precise genetic tests about their dispositions for severe diseases. They are afraid that a bad test result could render them desperate due to anticipatory feelings. We show that partially revealing tests are typically optimal when anticipatory utility interacts with an instrumental need for information. The same result emerges when patients rely on probability weighting. Optimal tests provide only two signals, which renders them easily implementable. While the good signal is typically precise, the bad one remains coarse. This way, patients have a substantial chance to learn that they are free of the genetic risk in question. Yet even if the test outcome is bad, they do not end in a situation without hope.

JEL Classification: D81, D82

Keywords: Test Design, Revelation of Information, Design of Beliefs, Medical Tests, Anticipatory Utility, Huntington's Disease

---

*We would like to thank Konstanze Albrecht, Doug Bernheim, Stefano DellaVigna, C. Engel, Péter Esö, Armin Falk, Russell Golman, W. Henn, Florian Herold, Rudi Kerschbamer, Emir Kamenica, Botond Kőszegi, Thomas Kruse, George Loewenstein, Matthew Rabin, Frank Rosar, J. Schneider, Andrei Shleifer, Alp Simsek, Josh Tasoff as well as seminar participants at SITE 2012, Bavarian Micro Day 2012, EBEM 2013, ESEM 2013 and at Bonn and Innsbruck Universities. Financial support of the DFG through, respectively, Hausdorff Center for Mathematics and SFB TR 15 is gratefully acknowledged. Nikolaus Schweizer, Department of Econometrics and OR, Tilburg University, email: n.f.f.schweizer@uvt.nl. Nora Szech, Department of Economics, Karlsruhe Institute of Technology, Berlin Social Science Center, and CESifo, email: nora.szech@kit.edu.

1 Introduction

It is one of the most elementary principles of decision theory that agents prefer to have as much information as possible before making a decision. More information allows agents to fine-tune their decisions. For example, planning the future becomes easier when knowing the challenges that lie ahead.

The following thought experiment may illustrate why there are nevertheless contexts in which information is not necessarily desirable: Imagine there was a test that could determine whether you survive the next t years or not. Set t to a relevant value, e.g., about half the time you expect to survive from now on. Assume that for some reason you are entirely confident about the accuracy of the test. Would you want to get this information? Contrary to the reasoning in the first paragraph, this is a question many people find difficult to answer.

While most people are currently not confronted with such a fundamental testing decision, people under risk of having inherited Huntington's disease are. Huntington's disease is a severe genetic disorder which breaks out around the age of 40. As more and more cells get damaged by the disease, both mental and physical health deteriorate. After some years, patients end up in dementia and disability, needing full-time care. Patients die 20 years younger than other people on average. There is no cure for Huntington's disease. Children of patients have a 50% risk of having inherited the disease (provided that exactly one parent has it). Consequently, the risk for grandchildren is 25%. Since the 1980s, a genetic test is available which allows laboratories to almost perfectly determine if a person will eventually get the disease.

People under risk often find it difficult to decide whether to take the test or not. There are books solely dedicated to this decision.[1] Many people postpone the decision, they wait for years before eventually taking the test. The problem of testing for Huntington's disease may seem like a – disturbingly severe – minority problem. Yet it is likely that this type of problem will become more wide-spread as research into human genetics progresses and more and more genetic dispositions become detectable. We will employ this testing decision as our running example in the following.

[1] See, e.g., Baréma (2005).

Our paper studies test design in this context. In our model, we blend risk preferences regarding anticipated outcomes with instrumental information. Optimal tests turn out to consist of two signals as follows. The good signal is precise and proves that the patient will stay healthy. The bad signal is coarse, implying that the patient has to correct his Bayesian belief of staying healthy downwards. Yet there is still hope that the disease will not break out as the bad signal is a pooling signal. Using the term "positive" as in "HIV-positive," the test thus does not yield false negatives, but allows for false positives. The optimality of this test structure is robust to the absence of instrumental value of information if the patient's anticipatory utility is partly concave and convex.[2] It also remains optimal if the patient applies probability weighting as in Kahneman and Tversky (1979).[3]

This paper presents a simple but robust model that aims at capturing why tests that provide life-changing outcomes can be challenging. To this end, we blend risk preferences regarding anticipated outcomes with instrumental information. We study the design of optimal tests in this setup. Optimal tests often turn out to be coarse. They either provide a perfectly informative signal in the good domain, then the genetic risk is absent, or they provide a coarse bad signal, which means that the patient has to correct his hope to stay healthy downwards. In the latter case, there still remains justified reason to hope that the disease will not break out as the coarse signal emerges from pooling. This differentiates the coarse test from the precise one.

This type of test is easily illustrated in terms of our thought experiment from the beginning. Imagine that with a probability of 50% you will survive the next t years. Consider a test that provides only two outcomes: If you will live for more than these t years, the test reveals this with a probability of, say, 30%. In all other cases, you receive a pooling signal which implies that you have to adjust your life expectancy slightly downwards. Thus, taking the test offers the possibility of getting good news while you never receive information that you will die within the next t years for sure. Taking such a test may feel less scary than a precise one.

The medical literature has discussed the careful use of precise tests extensively[4] but has not

[2] See Section 4 as well as Baucells and Belezza (2017).

[3] See Section 5.

[4] For example, basic fertility tests, though cheap, are recommended only to couples who have unsuccessfully tried for one year to become pregnant, see the current guidelines of the CDC or the British NHS. Recently, the

looked into the possibility of letting patients choose between precise and coarse tests. While new in medical testing, randomized mechanisms are well-established in a variety of other contexts. Examples range from complex random procedures for determining start configurations in sports contests (such as in the soccer world cup) to randomized pricing in the airline industry.[5] We emphasize in this paper that the technical progress in the design of information structures allows for better test design when it comes to crucial, potentially life-changing tests as in medical contexts as well.

The key idea behind our model is that an agent's utility at a given point in time is influenced not only by his current situation but also by expected future prospects. This is the anticipatory utility approach put forward by Loewenstein (1987), see below for more references. For a very simple example of anticipatory utility, people look forward to holidays in Hawaii and this may lift their spirits even months before the journey begins. Notably, what influences their utility *now* is not how those holidays will actually turn out to be, but how they expect them to be. This idea is subtly yet crucially different from the classical assumption that an agent only takes into account the (discounted) utility he enjoys at a later point in time when making a decision.

Agreeing to receive a piece of information is, so to say, equivalent to entering a gamble over anticipated utility outcomes which – by Bayesian rationality – leaves the status quo unchanged in expectation. Thus, other factors aside, an agent who is risk-averse with respect to anticipated utility will never want to receive any information about the future – while an agent who is "risk-loving", i.e. eager to learn about the future, would opt for precise information. The risk aversion in our model is hence analogous to the standard concept of risk aversion with the only difference being that it applies to anticipated outcomes instead of realized physical outcomes.

In addition to anticipatory utility, our model incorporates utility losses due to insufficient planning. Agents with better information make better decisions, e.g., better-suited career or family plans. Under risk aversion regarding anticipated payoffs, this leads to a trade-off. Getting more

PSA test (an indicator for prostate cancer) was criticized heavily for being overused on patients, see, e.g., Walter et al. (2006).

[5] See also Kamenica and Gentzkow (2011) for a recent contribution on strategic randomization in information transmission as applicable, e.g., in litigation.

information allows to make better plans for the future, but it also increases the risk of obtaining bad information that will lower anticipatory utility.

Our model and analysis can easily be augmented to include other behavioral aspects. Anticipatory utility has been identified as a plausible factor in decisions about receiving crucial information about the future. There may be alternative or additional reasons why patients shy away from medical tests. Likewise, avoiding costs from less than optimal plans for the future need not be the only argument in favor of perfectly revelatory tests.[6] We design optimal tests when there are conflicting forces at work. As a concrete example, we adapt the analysis to a setting where patients use probability weighting to evaluate the likelihood of different health outcomes. The difference between being healthy with a probability of 90% or 100% may be perceived as much greater than the difference between 50% and 60%. This is the famous "underweighting of high probabilities" pointed out by Kahneman and Tversky (1979). Even in the absence of instrumental information, probability weighting leads to the same structure of optimal tests as before. Probability weighting is thus a second, independent factor in favor of tests that either deliver clear-cut good news, or a coarse bad signals.

1.1 Related Literature

The idea of anticipatory utility goes back to Bentham.[7] Contributions such as Loewenstein (1987), Caplin and Leahy (2001), Brunnermeier and Parker (2005), Epstein (2008), Kadane, Schervish and Seidenfeld (2008), Golman and Loewenstein (2012) and Baucells and Belezza (2017) have developed concepts of anticipatory utility in behavioral economics and decision theory.[8] Building on this research, Caplin and Leahy (2004), Caplin and Eliaz (2003), Kőszegi (2003, 2006) and Oster, Dorsey and Shoulson (2013) study information transmission in doctor-patient relationships. The main distinction between our work and most of these contributions is that we focus on the design of optimal tests that can be partially revelatory.

Caplin and Leahy (2004) study testing decisions under anticipatory utility yet in the absence of

[6] Another reason could be curiosity, which, as we will see, can be covered by the analysis as well.

[7] See Read (2007) and Morewedge (2016) for overviews.

[8] For empirical tests, see e.g. Chan and Mukhopadhyay (2010), Ganguly and Tasoff (2016) and Huck, Szech, and Wenner (2015).

instrumental information. Kőszegi (2003, 2006) blend anticipatory utility with costs of suboptimal decisions. Kőszegi (2003) focuses on patients' preferences with regard to perfectly revelatory tests. Kőszegi (2006) studies the exchange of information between doctor and patient in a cheap-talk game in which the doctor is severely limited in his power to commit to truthfulness. The patient has to choose between taking a therapy or not. For the doctor, as he cares about the patient's well-being, this creates an incentive to downplay the severity of the patient's illness (as long as the patient still takes the therapy).[9] Yet the patient understands this and therefore the doctor can only credibly release rough signals about the health status of the patient.[10] In a sense complementary to this analysis, we focus on the case where commitment is possible as, e.g., in genetic testing, as hard information can be generated in this case and the doctor does not receive more information than the patient himself. See also the discussion at the end of Section 2. Finally, in an empirical study, Oster, Dorsey and Shoulson (2013) show that anticipatory utility can well explain observed decisions for and against taking the perfectly revelatory test for Huntington's disease.

To our knowledge, Caplin and Eliaz (2003) is the only other paper which considers optimal test design under anticipatory utility. The authors focus on tests for HIV. They suggest partially revelatory certificates as a way to motivate agents with anticipatory utility to get tested at all. Caplin and Eliaz identify a testing procedure that yields an "infection-free" equilibrium such that HIV is no longer transmitted to healthy people. The test is thus designed with a different intent, namely protecting healthy people. The well-being of an individual patient is a potential constraint that does not allow the test-designer to implement the first best solution which would be a perfectly revelatory test. Thus there is a conflict of interest between test-designer and potentially infected patients. This is different from our analysis, where doctor's and patient's interests are perfectly aligned. The goal is to identify the optimal test according to an individual patient's needs.

Eliaz and Spiegler (2006) argue that models of anticipatory utility cannot capture the following

[9] The doctor would thus release messages like "It looks basically fine but we may nevertheless do the therapy to be on the safe side..."

[10] Kőszegi (2006) also briefly considers the case where the doctor can commit on truthful revelation. Yet under his assumptions on preferences and costs, this always leads to full revelation.

stylized fact: Patients with a high prior expectation of being healthy are more information-seeking than patients with a low prior. As we demonstrate in Sections 4 and 5, their "impossibility result" does not apply if tests including a coarse signal turn out optimal. We find that the optimal test is often non-revealing for small priors, yet partially revealing for large priors. We discuss the relation to Eliaz and Spiegler (2006) in detail at the end of Section 4.

There has been some debate about revealed-preference foundations for anticipatory utility (Eliaz and Spiegler, 2006, Epstein, 2008). Furthermore, it may be difficult for a doctor to infer a patient's exact preferences for information. Our contribution is, in a sense, orthogonal to this discussion. We argue that the same small family of tests emerges as optimal under a broad class of preferences. Thus, offering the patient some tests from that family can be a good idea even without knowledge of his exact preferences.

From a technical point of view, our paper is related to works on strategic conflict in information transmission: Rosar (2014) also considers test design. The characterization of optimal signals in Kamenica and Gentzkow (2011) crucially relies on a classical result from geometric moment theory (Kemperman, 1968) which is also the key in our derivation of optimal tests. These papers focus on problems caused by strategic interaction between economic agents. We consider problems caused by the need to control one's own expectations. In our model, there is no conflict of interest when it comes to information transmission.

2 The Basic Model

Consider the following game between a receiver of information ("the patient") and a revealer of information ("the doctor"). Doctor and patient share the goal of maximizing the patient's utility. They also have the same information about the patient's preferences and the ex-ante situation. There is an initially unknown state of the world X which takes the values 1 and 0 with commonly known probabilities p and $1-p$. In applications such as genetic testing, such priors only hinge on which relatives of a patient are known to have the disease. Throughout, $X=1$ denotes the preferred outcome ("the patient is healthy"). $X=0$ denotes the unfavorable outcome ("the patient has a severe genetic mutation and will become ill"). The timing of decisions is as follows.

(i) The doctor designs a test for the patient. A test is the specification of a joint distribution (S, X) such that the marginal distribution of X is Bernoulli with mean p, and S is a random variable correlated with X. As it turns out,[11] the optimal test necessarily takes only two values. Hence, the doctor's search space ultimately reduces to a 2×2 matrix specifying the joint distribution of (S, X).

(ii) The patient learns the joint distribution of S and X. He decides whether he wants to take the test and observe the realization of S, or not.

(iii) The patient forms a Bayesian posterior belief B about the distribution of X. If he has taken the test, his belief to be healthy adjusts, $B = P[X = 1|S]$. If he opted against the test, B remains the prior belief, $B = p$.

(iv) The patient settles on a life plan, modeled as choosing a value $y \in [0, 1]$.

The shared goal is to optimize the patient's expected realized utility. Realized utility consists of three terms.

$$U_c(X) + \theta U_a(E[X|B]) + (1 - \theta)U_p(X, y).$$

Note that the middle term involving U_a hinges on the patient's posterior expectations of being healthy. The remaining two terms depend on the actual realization of X. Here, the term $U_c(X)$ captures the classical, "physical" utility from being healthy or ill.

In the middle term, $U_a(E[X|B])$ captures anticipatory utility. It is a function of the patient's Bayesian posterior expectation of X.[12] A patient may feel miserable knowing that he will become ill. Anticipating this, he may want to avoid a too revelatory test about his health condition. The patient is thus averse to fluctuations in $E[X|B]$. Accordingly, we assume U_a as increasing, continuous and concave.[13]

[11]See Proposition 1 below.

[12]A credible testing procedure will have to obey the rules of Bayesian statistics. Resulting posteriors are communicated explicitly to the patient. This rules out misperceptions of probabilities. For potential effects of probability weighting, see Section 5.

[13]This assumption is in line with classical theories of anticipatory utility as in Loewenstein (1987). As argued in Baucells and Belezza (2017), there are also empirical indications that U_a is increasing and switches from being concave to being convex at some point. This captures that information avoidance may be limited to unfavorable realizations. As discussed in Section 4, our main results on optimal tests also cover this setting.

The planning utility term U_p models how the patient can influence his condition by a careful choice of life plan y. This terms captures the instrumental value of information. $U_p : \{0,1\} \times [0,1] \to \mathbb{R}$ has the property that for fixed $x \in \{0,1\}$, $U_p(x,y)$ is continuous in y and takes its unique maximum in $U_p(x,x) = 0$. Planning utility is thus maximal when the patient knows the state of the world X and can choose the best-suited life plan $y = X$. If the patient does not know X, he cannot adjust his plans optimally to his (future) health condition. This leads to a (typically negative) utility out of suboptimal planning, $U_p(X,y)$. For example, the patient may want to opt for a different career plan, travel more or take more leisure time, take better care of his savings, or buy a home close to his family instead of moving far away or even abroad if he knows he is going to become ill eventually. Note that U_p only captures deviations from the overall utility that can be achieved for a given realization of the health outcome X. In contrast, fluctuations in the overall utility level depending on the health outcome are covered by the classical term $U_c(X)$.

$\theta \in [0,1]$ is a parameter that captures how important anticipations are compared to choosing a good life plan. We will later vary θ in order to investigate the interplay of the two terms.

We think of U_c, U_a, and U_p as aggregates over all future time periods, i.e., discounted sums of future physical utilities, future anticipations, and future utility from having chosen a life plan y which is – ex post – suboptimal. Likewise, the choice of life plan y should be understood as an aggregate over many decisions (occupational choice, investment and saving plans, etc.). If we think of X as a genetic indicator of whether a disease will eventually break out, the patient inevitably observes X in the far future. Accordingly, the life plan y only captures decisions made before that point in time.

Two simplifications are immediate. First, the physical utility term $U_c(X)$ is unaffected by the testing decision. It thus does not play any role in the later analysis. We will therefore omit this term without loss of generality. Second, the doctor will always propose the best possible test for the patient. This may be a test containing pure noise from which the patient cannot learn anything. We can thus assume without loss of generality that the patient accepts the test in stage (ii).

We assume that the doctor can offer the test the patient likes best. For instance, he may send instructions for generating the test to a medical laboratory. With regard to genetic testing, many noisy signals can be created by mixing blood samples of different patients and just testing the mixed sample for the genetic mutation of interest.[14] Assume the blood of two people at risk is mixed and then tested with a precise test: If the mixed blood sample is clean, both patients are free of the genetic mutation. If the mixed blood sample contains the mutation, either one of the patients carries the mutation, or both do. Of course, another way to generate noisy signals is via computerized, anonymous processes.

3 Optimization

In this section, we derive the general structure of optimal tests. We apply backward induction. First, we optimize the patient's planning decision in reaction to a given posterior belief induced by the test. Taking into account how the patient will react, we then optimize over posterior beliefs. The resulting pairs of optimal tests and life plans can be interpreted as Perfect Bayesian Equilibrium (PBE) of the doctor-patient game. Yet as the interests of patient and doctor are aligned, it is in fact irrelevant whether the patient or the doctor designs the test. We can thus also interpret the solution as the result of a two-stage optimization executed by the patient.

Proposition 1 and Lemma 2 characterize, respectively, the beliefs induced by an optimal test and the optimal test itself. Proposition 2 describes how the optimal test becomes more revelatory if the costs of making wrong decisions become more important compared to the anticipatory effects.

Let us first consider the patient's choice of y given that B has taken the realization $B = b$. Ignoring terms that are independent of y, the patient's problem to choose a good life plan y is given by

$$\max_y u_p(b, y) \text{ where } u_p(b, y) = b\, U_p(1, y) + (1 - b)\, U_p(0, y).$$

[14]In other medical contexts, e.g., to ensure safety of blood donations, testing mixed samples is a known procedure for reducing costs, see Section 6 for more discussion.

Since $u_p(b, y)$ is continuous in $y \in [0, 1]$, an optimal choice of y exists for all b. We denote it by $y^*(b)$. The planning utility given that the patient behaves optimally is thus

$$u_p^*(b) = u_p(b, y^*(b)).$$

Our first result shows that u_p^* is convex in b.

Lemma 1. *The function $u_p^*(b)$ is non-positive, continuous and convex in $b \in [0, 1]$ with $u_p^*(0) = u_p^*(0) = 0$.*

All proofs are in the appendix. We now turn to the doctor's task of designing the optimal test for the patient. We take an indirect approach. First, we determine the optimal belief B^*. Then we construct a test that induces this belief. To this end, denote by $\mathcal{B}$ the set of random variables valued in $[0, 1]$ that have mean p. By Bayesian consistency, the doctor cannot induce any belief outside of $\mathcal{B}$, as the prior needs to be preserved in expectation.[15] The set $\mathcal{B}$ thus encodes all possible tests, including the special cases of perfect revelation ($B \in \{0, 1\}$) and no revelation ($B = p$ a.s.).

The doctor aims at maximizing the patient's expected utility, assuming the patient chooses the conditionally optimal life plan y^* based on the test result:

$$\max_{B \in \mathcal{B}} E[V(B)] \text{ where } V(b) = \theta U_a(b) + (1 - \theta) u_p^*(b). \tag{1}$$

Here, we have ignored the term $E[U_c(X)]$ since, by the law of iterated expectations, it does not affect the maximization problem. Furthermore, we have used that $E[X|B] = B$ and thus $U_a(E[X|B]) = U_a(B)$. By assumption, U_a is concave. Moreover, Lemma 1 has shown that u_p^* is convex. Thus, for $\theta \in (0, 1)$, the function V is generally continuous but neither convex nor concave. This stems from the conflict that lies at the heart of the problem: The utility from choosing a good life plan, $E[u_p^*(B)]$, demands resolution of uncertainty. Yet the anticipatory utility term, $E[U_a(B)]$, suggests to avoid information.

[15] It can be shown that the doctor can induce any $B \in \mathcal{B}$, see Shmaya and Yariv (2009). Since we will first determine the optimum $B^* \in \mathcal{B}$ and then implement it directly, this type of result is not needed here.

U_a does not need to be globally concave for this conflict to arise. As soon as V is non-convex, the optimal test should not be fully revealing for some priors p. Similarly, our analysis is robust to further psychological factors such as anxiety, curiosity, fear, etc. The sole property of V that is used in the following is that it is a continuous function.[16] For instance, we could add a term $\gamma F(b)$ modeling curiosity to the function V. In order to capture that a more informative signal satisfies the patient's curiosity better, we could assume that F is strictly convex. This would not require any changes to our analysis (and we could conclude that for sufficiently large γ the incentives for receiving as much information as possible become dominant).[17]

The optimization problem (1) is a classical problem in geometric moment theory which was solved independently by various authors in the 1950s. We refer to Kemperman (1968) for an overview of the earlier literature. To our knowledge, Richter (1957) contains the first published statement of a result which immediately implies Proposition 1 below.[18] For ease of reference, we provide a short and non-technical exposition of how to solve (1) which is given in the proof of Proposition 1.

The key observation is that the patient's utility from the optimal test is given by $\overline{V}(p)$ where $\overline{V}$ is the smallest concave function greater than or equal to V. Moreover, the optimal test can be read off from the graph of $\overline{V}$ as is depicted in Figure 1.

For illustration, consider a test inducing a belief B that takes only the two values $d_l < p < d_h$. The patient's utility from this test can be found graphically by connecting the points $(d_l, V(d_l))$ and $(d_h, V(d_h))$ and evaluating the value of the resulting line segment at p. Since $\overline{V}$ can be characterized as the supremum over all line segments which connect two points in the graph of V, $\overline{V}(p)$ is exactly what the optimal test can achieve. The proof of Proposition 1 demonstrates

[16] This continuity is a convenient technical assumption since it implies that V attains intermediate values, maxima and minima. It can be relaxed at the expense of more complicated statements of the results.

[17] Formally, a similar trade-off between concavity and convexity also lies at the heart of Kamenica and Gentzkow (2011)'s results on partial revelation in strategic information transmission. In their setting, one agent, the sender, can observe the state of the world. His utility depends on the decision made by another agent, the receiver, who uses the information he receives to maximize his own utility. When preferences are aligned, full revelation is optimal in their setting. In contrast, in our analysis it turns out that full revelation is often dominated by less revelatory information structures.

[18] For a broader perspective on moment problems, generalized Chebychev inequalities and applications in decision analysis, see Smith (1995). The result is also a key ingredient in Kamenica and Gentzkow's (2011) analysis of strategic information transmission. Earlier applications of similar techniques in the context of strategic information transmission are found in Aumann and Maschler (1995).

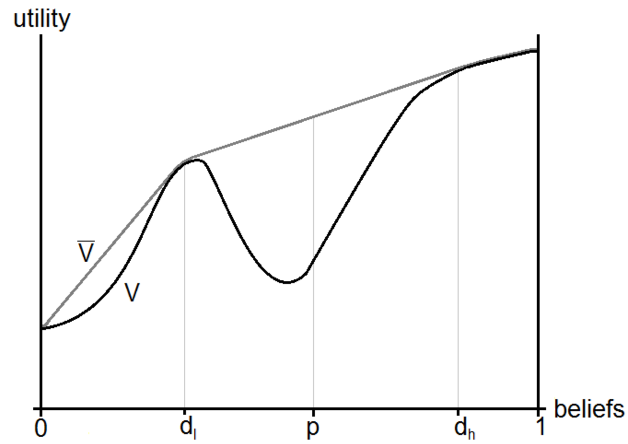

Figure 1: Construction of $\overline{V}$

this point in more detail. It also shows that beliefs B which take more than two values cannot achieve more than $\overline{V}(p)$.

Proposition 1. *Denote by $\overline{V}$ the smallest concave function with $\overline{V}(b) \geq V(b)$ for all $b \in [0,1]$. Then a solution $B^* \in \mathcal{B}$ to (1) is given as follows:*

(i) If $\overline{V}(p) = V(p)$ then $B^* = p$ with probability 1.

(ii) If $\overline{V}(p) > V(p)$ denote by $I = (b_l, b_h) \subset [0,1]$ the largest open interval with $p \in I$ and $\overline{V}(b) > V(b)$ for all $b \in I$. Then B^* takes values b_h and b_l with probabilities

$$p_h = \frac{p - b_l}{b_h - b_l} \text{ and } p_l = 1 - p_h.$$

In both cases, $E[V(B^*)] = \overline{V}(p)$.

Existence of $\overline{V}$ is ensured since the convex hull of the graph of V exists and $\overline{V}$ is the upper contour of that convex hull. It is easy to check that B^* is unique if there are no subintervals of $[0,1]$ on which V is linear.

To get some more intuition for the objects in the proposition, consider the case of $\theta = 0$, i.e., the case of a patient who only cares about early resolution of uncertainty. Then V is convex and accordingly, $\overline{V}$ is given by the straight line connecting $(0, V(0))$ and $(1, V(1))$. In that case, $\overline{V}(b) > V(b)$ for all $b \in (0,1)$ and the proposition implies that B^* takes values 0 and 1 with probabilities $1-p$ and p. Thus the patient perfectly learns from the test whether $X = 0$ or $X = 1$. In the case where $\theta = 1$, i.e., for a patient whose interests are dominated by anticipatory utility, V is concave and thus $\overline{V} \equiv V$. Accordingly, we are in case (i) of the proposition and the optimal belief B^* coincides with the prior p. Hence the optimal test does not reveal anything. In the case where V and $\overline{V}$ coincide on some interval, it depends on the prior p whether the optimal test should reveal something or not. As we show in Section 4, one concrete class of examples where optimal tests are partially revealing in general is as follows: U_a satisfies decreasing absolute risk aversion while U_p corresponds to a quadratic distance between the ex-post optimal and the actual life plan.

Proposition 1 characterizes the structure of optimal beliefs. In particular, it shows that optimal beliefs B^* lie in the subset $\mathcal{B}_2 \subset \mathcal{B}$. Here, $\mathcal{B}_2$ is defined as the set of random variables on $[0,1]$ which have mean p and which take only two values b_l and b_h where $b_l \leq b_h$. Thus, to derive the optimal signals, it suffices to show that for any $B \in \mathcal{B}_2$ there exists a signal S which induces B. This is the result of the following lemma.

Lemma 2. *Fix $0 \leq b_l < p < b_h \leq 1$ and consider the random variable S with values in* {"Good", "Bad"} *that is generated upon observing X as follows:*
If $X = 1$ then

$$S = \begin{cases} \text{“Good”} & \textit{with probability } \alpha \\ \text{“Bad”} & \textit{with probability } 1-\alpha. \end{cases}$$

If $X = 0$ then

$$S = \begin{cases} \text{“Good”} & \textit{with probability } \beta \\ \text{“Bad”} & \textit{with probability } 1-\beta, \end{cases}$$

where $\alpha, \beta \in [0,1]$ are given by

$$\alpha = \frac{b_h}{p}\frac{p-b_l}{b_h-b_l} \quad \textit{and } \beta = \frac{1-b_h}{1-p}\frac{p-b_l}{b_h-b_l}.$$

The resulting belief $B = P[X = 1|S]$ only takes values in $\{b_l, b_h\}$ and $E[B] = p$.

Here, $S =$ "Good" is better news than $S =$ "Bad" since it induces the higher posterior probability b_h of the good state of the world $X = 1$. It is straightforward to rewrite the test of Lemma 2 in a way that X only needs to be observed with some probability.

We close this section with some qualitative results on optimal tests. The first result confirms the intuition that smaller values of θ – representing a higher significance of the cost term – lead to more precise tests.

Proposition 2. *Fix $p \in (0,1)$ and $\theta > \theta'$. Denote by $\{b_l, b_h\}$ and $\{b'_l, b'_h\}$ the values taken by the optimal belief under, respectively, θ and θ'. Then $b_l \geq b'_l$ and $b_h \leq b'_h$. Thus the optimal test under θ' leads to beliefs which are closer to knowledge of X than the optimal test under θ.*

The next result further illustrates the structure of optimal tests and states the following: Consider only tests which take two values and fix the lower of the induced beliefs d_l to a value which is less informative than optimal, $d_l \in (b_l, p)$. What is the optimal induced upper belief d_h^*? Proposition 3 shows that $d_h^* \in (p, b_h]$, implying that if a test is less informative than optimal in one direction, it is best to leave it less informative than optimal in the other direction, too.

Proposition 3. *Define the prior p and the values of an optimal belief $\{b_l, b_h\}$ as above. Assume that $b_l < p < b_h$ and fix some $d_l \in (b_l, p)$. For $d_h \in (p, 1)$, denote by $D(d_l, d_h) \in \mathcal{B}$ the random variable with mean p which takes only values d_l and d_h. Assume there exists d_h such that $E[V(D(d_l, d_h))] > E[V(p)]$ so that some beliefs $D(d_l, d_h)$ are better than no information. Then, if d_h^* is a solution to*

$$\max_{d_h} E[V(D(d_l, d_h))],$$

it must hold that $d_h^* \leq b_h$.

We have considered the case where d_l is fixed and d_h is variable. The argument for the opposite case is analogous.

4 Accuracy on Good News

Even though the optimal test is simple in the sense that it just provides two potential results, there may still be challenges when designing it. In this section, we restrict the functions $U_a(\cdot)$

and $U_p(\cdot,\cdot)$ a bit more. The quality of a life plan will hinge on how far away it is from the ex post optimal one, and the patient will be specifically scared about receiving very bad news. Designing the optimal test then reduces to determining one single parameter. The optimal test structure becomes as follows: The test may perfectly reveal the good state of the world, but it never perfectly reveals the bad state. In other words, there are no false signals of disease-freeness, while false positives occur. A patient thus either learns that he remains healthy for sure, or he receives a pooling signal. In the latter case, his belief of staying healthy is corrected downwards, but not to zero.[19] In the terminology of Lemma 2, the optimal test is characterized by $\alpha \in (0,1)$ and $\beta = 0$. Such a test structure emerges for example if V is concave on pessimistic beliefs and convex on optimistic ones as we will see in the following.

Assumption 1. *Let V be continuously differentiable and assume there exists a point $b_c \in (0,1)$ such that $V(b)$ is strictly concave on $[0,b_c]$ and strictly convex on $[b_c,1]$.*

The idea is that for pessimistic beliefs about staying healthy, the anticipatory utility term is dominant, while the cost term dominates for optimistic beliefs. Such situations occur if costs of making suboptimal plans hinge on the distance to the ex-post optimal plan, while anticipatory utility is more sensitive to small changes in beliefs near the undesirable diagnosis $X = 0$. If patients are specifically scared of ending in a situation of no or very little hope, this assumption should be fulfilled. The following example provides concrete functional assumptions on U_a and U_p which describe such a setting. The calculations are in the appendix.

Example 1. *Suppose U_a is three times continuously differentiable with $U_a''' > 0$. For $x \in \{0,1\}$, the function $U_p(x,y)$ is given by $U_p(x,y) = -(x-y)^2$. In this case, there exist thresholds $\theta_l \leq \theta_h$ in $[0,1]$ such that V is strictly concave for $\theta \geq \theta_h$ and strictly convex for $\theta \leq \theta_l$. For all $\theta \in (\theta_l,\theta_h)$, there exists a point $b_c \in (0,1)$ such that $V(b)$ is strictly concave on $[0,b_c]$ and strictly convex on $[b_c,1]$.*

In the example, either anticipatory utility or planning concerns dominate for extreme values of θ. This leads to perfectly revelatory or non-revelatory optimal tests. For intermediate values of θ

[19] A similar class of tests was found optimal in Rosar (2014) in a model of strategic conflicts in information transmission. Caplin and Eliaz (2003) choose this type of test for implementing an "infection-free" equilibrium in their model of testing for AIDS.

we are in the setting of Assumption 1. Anticipatory utility dominates at pessimistic beliefs while planning utility dominates for optimistic beliefs. The idea behind the example is the following. Regarding anticipations, learning a bit more feels particularly risky if very bad outcomes are possible.[20] With regard to the utility losses from an unsuitable life-plan, it is only the distance to the ex-post optimal plan that matters.

A different type of setting which satisfies Assumption 1 arises when the instrumental value of information is absent (or negligible), $\theta = 1$, but when the function U_a itself is first concave and then convex. This corresponds to a patient who is reluctant to receive bad news but curious about good news. See Baucells and Belezza (2017) for some foundations for this shape of U_a. In Section 5 below, we demonstrate that the interplay of concave U_a with certain forms of probability weighting can effectively lead to similar situations.

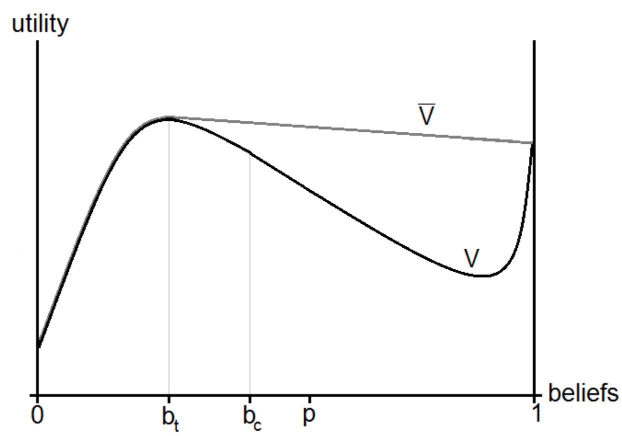

Figure 2: Construction of $\overline{V}$ under Assumption 1

In the concave-convex setting of Assumption 1, the function $\overline{V}$ has a particularly simple structure depicted in Figure 2. There exists $b_t \geq 0$ such that $\overline{V}$ is identical to V for $b \leq b_t$. Over the interval $(b_t, 1)$ it is given by the straight line connecting $(b_t, V(b_t))$ and $(1, V(1))$. Accordingly, the optimal test induces posterior beliefs b_t or 1. The case $b_t = 0$ corresponds to the case where

[20]Recall that U_a is concave and thus U_a'' being increasing means that $U_a''(b)$ is closer to zero for larger b. The assumption of an increasing second derivative of U_a, $U_a''' > 0$, was coined "prudence" by Kimball (1990). It is a necessary condition for decreasing absolute risk aversion and thus satisfied by many of the standard utility functions.

the straight line connecting $(0, V(0))$ and $(1, V(1))$ dominates the graph of V for all b. In this case, full revelation is optimal at all priors. Otherwise, as illustrated in the figure, b_t is the unique point at which the tangent to V passes through $(1, V(1))$.

Proposition 4. *Under Assumption 1, there exists $b_t \in [0,1]$ such that the optimal test is as follows.*

(i) If $p \leq b_t$, the optimal test is perfectly non-revealing, e.g., $\alpha = \beta = 0$.

(ii) If $p > b_t$, the optimal belief B^* takes only values $b_l = b_t$ and $b_h = 1$. The resulting optimal test is given by

$$\alpha = \frac{1}{p}\frac{p - b_t}{1 - b_t} \quad \textit{and} \quad \beta = 0.$$

Here, α and β are as defined in Lemma 2. Thus, for $b_t > 0$, the optimal test sometimes reveals $X = 1$ but never $X = 0$.

The analysis in this section shows that simple binary tests with $\alpha \in [0,1]$ and $\beta = 0$ may be promising candidates to include into menus of tests. If a doctor wishes to propose some test options to a patient – in addition to the perfectly revelatory and perfectly non-revelatory tests represented by $\alpha \in \{0,1\}$ and $\beta = 0$ – it might be a good starting point to include a discretization of the range of α, e.g. the three tests corresponding to $\alpha \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$, and $\beta = 0$.

In the case $b_t > 0$ of Proposition 4, we see that patients with a small prior probability of the favorable outcome[21] refuse any further information. Patients with a larger prior will instead like a partially revealing test best.[22] The demand for information thus depends on the prior.

In contrast, Eliaz and Spiegler (2006) have argued that anticipatory utility cannot explain this intuitive type of preference reversal (see their Example 2). Their critique is based on the following result (their Proposition 2): Suppose the fully revealing test is either the best possible or the worst possible test for p close to 0 or for p close to 1. Then the fully revealing test is either

[21] These people are in a sad kind of lottery-like situation as they only have a small chance of the good outcome.

[22] Moreover, under Assumption 1, patients with a sufficiently large prior typically prefer complete information over no information. This holds whenever the straight line connecting $(0, V(0))$ and $(1, V(1))$ is greater than V near 1.

best or worst for all $p \in [0,1]$.[23] Yet note that this claim does not stand in conflict with our results.[24] For $b_t > 0$, the fully revealing test is never optimal. For all priors, full revelation is dominated either by no revelation or by partial revelation. With a similar argument, one finds that full revelation is never the worst either. Instead, a partially revealing test turns out as the worst for small priors, and a perfectly non-revealing test is worst for large priors.[25] Thus, under Assumption 1, the premise of the analysis of Eliaz and Spiegler (2006) that full revelation is either best or worst is typically not fulfilled.

5 Biased Perceptions of Probabilities

The previous section showed that optimal tests are more accurate on good than on bad news whenever Assumption 1 is fulfilled, i.e., whenever the function V is concave up to some point and convex from there on. This section demonstrates that such situations also arise if the patient relies on probability weighting – even if there is no instrumental value of information. Probability weighting, i.e. a biased perception of probabilities, is thus an alternative argument in favor of the structure of optimal tests identified in Proposition 4.

In the following, let us set the planning component of overall utility to zero, i.e., $\theta = 1$. Unlike before, we assume that the patient's anticipatory utility does not depend on the posterior belief B directly, but rather on a weighted version $w(B)$ of that belief. Such probability weighting has been discussed extensively in the behavioral literature, mostly in the context of prospect theory and rank-dependent utility.[26] Following that literature, we assume that the probability weighting function w is an increasing function with $w(0) = 0$ and $w(1) = 1$ so that probabilities of certain events are evaluated correctly. Furthermore, the typical form of w is an inverse S-shape: w grows quickly near 0 and near 1. The patient thus perceives differences between

[23]This can be seen as follows: Proposition 1 implies that for any prior full revelation can only be best or worst when the straight line connecting $(0, V(0))$ and $(1, V(1))$ lies above or below the graph of V for all b. This a global property which is prior-independent.

[24]Epstein (2008) provides an alternative reconciliation using richer classes of preferences which may depend on the prior in more general ways than under classical, expectation-based anticipatory utility.

[25]The worst possible test can be read off from the smallest convex function below V in an analogous fashion to the optimal test. As seen in Figure 2, this smallest convex function is linear starting in $(0, V(0))$ until some point b_w in the concavity region where it is tangential to V. The worst possible test induces beliefs $b_l = 0$ or $b_h = b_w$ for $p < b_w$ while it is non-revealing for $p > b_w$.

[26]See Wakker (2010), especially Chapter 7, for an introduction.

intermediate probabilities, e.g. 40% and 70%, smaller than they really are. To such a patient, a coarse test that involves a clear-cut signal of disease-freeness may be specifically appealing. The chance of learning to be perfectly healthy outshines the potential risk of bad, yet still coarse, news. The following analysis confirms this intuition.

The doctor's test design problem now becomes

$$\max_{B \in \mathcal{B}} E[V(B)] \text{ where } V(b) = U_a(w(b)). \tag{2}$$

When optimizing over potential tests, the doctor takes an unbiased expectation over the patient's long-term well-being as affected by the respective test results.[27] [28]

We assume that the function w is three times continuously differentiable and strictly increasing with $w(0) = 0$ and $w(1) = 1$. Moreover, we assume that there exists $b_c \in (0, 1)$ such that w is strictly concave over $[0, b_c]$ and strictly convex over $[b_c, 1]$. The latter is the assumption of an inverse S-shape which is satisfied by most common smooth probability weighting function such as those of Tversky and Kahneman (1992) and Prelec (1998) over the commonly studied parameter regions. The following observation is immediate.

Corollary 1. *Suppose that the patient's anticipatory utility is linear in perceived beliefs, i.e., $V(b) = w(b)$. Then V satisfies Assumption 1.*

Thus, even for a patient who is neutral to variations in his biased probabilities, the optimal test is of the form described in Proposition 4 and thus involves coarse signals. This finding is in contrast to the case $V(p) = w(p) = p$ of unbiased perceptions of probabilities where any test is as good as any other since $E[V(B)] = p$ for any admissible B. We next return to the case where the patient is averse to variations in his biased beliefs so that U_a is concave. Proposition

[27] A patient will likely have to live many years with a test result, such that biases in risk perception could play important roles for the utility he derives from that result. Studies document that patients tend to estimate health risk from genetic disease with less bias than risk resulting from unhealthy behavior (Weinstein, 1984). This may be driven by an illusion of control rather that by an unrealistic optimism per se (compare McKenna, 1993). Yet also with regard to genetic risk, biases in perception have been documented in patients (e.g. Erblich et al., 2000).

[28] A patient may be in addition short-term biased regarding probabilities of test results. Yet such biases become irrelevant right after the test has been conducted. We therefore assume that the doctor optimizes the patient's welfare taking into account how the patient will feel about the test result in the decades to come but not how the patient may distort the probabilities of test outcomes right before the test is carried out.

5 formulates sufficient conditions such that the interplay of U_a and w induces a function V that is first concave and then convex and thus fulfills again Assumption 1.

Proposition 5. *Suppose U_a is twice continuously differentiable and DARA, i.e., $-U_a''(x)/U_a'(x)$ is decreasing. Denote by $h(y) = w^{-1}(y)$ the inverse of w and assume $-h''(y)/h'(y)$ is strictly increasing with $\lim_{y\uparrow 1} -h''(y)/h'(y) = \infty$. Then, the function $V(b) = U_a(w(b))$ satisfies Assumption 1.*

Intuitively, what needs to be ensured for the result is that the concavity of U_a dominates for small b while the convexity of w dominates for large b. The DARA assumption implies that U_a becomes increasingly less concave. The assumption on $h = w^{-1}$ goes in the opposite direction. Being the inverse of w, h is S-shaped and switches at b_c from convexity to concavity. The "IARA" assumption on probability weighting ensures that this switch is also a local property so that the function first becomes less and less convex and then more and more concave.

We next show that the probability weighting function proposed by Prelec (1998) satisfies the requirements of Proposition 5 over a wide range of parameters.

Proposition 6. *Let*

$$w(b) = \exp(-\kappa(-\log(b))^\rho)$$

where $\rho \in (0,1)$ and $\kappa < \rho^{-\rho}$. Then for $h(y) = w^{-1}(y)$ we have that $-h''(y)/h'(y)$ is increasing with $\lim_{y\uparrow 1} -h''(y)/h'(y) = \infty$.

$\rho \in (0,1)$ guarantees that w has an inverse S-shape rather than an S-shape. Moreover, as $\rho^{-\rho} > 1$, κ can take any value in $[0,1]$ and even larger values. The result thus covers the one-parameter case of Prelec's probability weighting ($\kappa = 1$) but also many cases with $\kappa > 1$, as long as κ does not become too large. For instance, Wakker (2010) proposes $\rho = 0.65$ and $\kappa = 1.05$ as plausible values. With $\rho = 0.65$, the restriction becomes $\kappa \leq \rho^{-\rho} \approx 1.32$ which clearly includes the case $\kappa = 1.05$.

The probability weighting function due to Tversky and Kahneman (1992)

$$w(b) = \frac{b^\rho}{(b^\rho + (1-b)^\rho)^{\frac{1}{\rho}}}$$

is analytically not as tractable as Prelec's but we can easily verify visually that its curvature satisfies our assumptions for common estimates of the parameter ρ. Here, we make use of the fact that instead of monotonicity of $-h''(y)/h'(y)$ we can equivalently study monotonicity of $w''(b)/(w'(b))^2$ as is shown in the proof of Proposition 5. As seen in Figure 4 in the appendix, the curve $w''(b)/(w'(b))^2$ is increasing and divergent at $b \uparrow 1$ in all cases. The values of ρ between 0.56 and 0.71 are taken from Table 1 of Bleichrodt (2001) who reports estimates from various studies.

6 Conclusion

Precise tests can be scary. Patients, for example, may shy away from tests that reveal whether they have a hereditary disease anchored in their genes or not, due to anticipatory feelings. We show that if there is an instrumental need for information, such as career or family planning, coarse test structures turn out optimal. Such tests typically avoid providing precise bad news. They give one of two signals – a precise good one, or a coarse bad one. The same test structure turns out optimal if patients rely on prominent forms of probability weighting with regard to their anticipatory well-being.

Creating such binary tests that involve coarse signals is easy. One way is to work with pooled samples. If several blood samples of different patients are mixed and only screened thereafter, the detection of a marker of disease does not imply that all people in the sample are concerned, but only that one certainly is. Testing pooled samples has been frequently used in other contexts. For example, in order to ensure safety of blood transmission while keeping costs manageable, pooled samples of donated blood are screened for infectious diseases, compare e.g. Stramer, Caglioti and Strong (2002). Another option is to work with computerized methods that involve randomization in the communication between the doctor and the laboratory.

Our paper does not suggest giving up perfectly revelatory testing. Rather, it argues that in many situations, providing a menu of tests including the precise test but also coarse ones, would be good. This way, patients can choose an information structure depending on their needs and feelings. So far, if patients are too scared to take the precise test, they are just left with the option to walk away and learn nothing.

With regard to Huntington's disease, some may think of more conventional economic explanations for avoiding the precise test. One may argue that conducting the test is costly, or worry that finding a health insurer becomes impossible if the test result is bad.[29] But neither of these explanations can fully capture what is going on. To see this, recall the thought experiment from the introduction about a reliable test which told you whether you would live for another t years. The test in the thought experiment is costless. Assume the test predicted only relatively early, sudden deaths. Then, finding a health insurer becomes easy in case of a bad test result. Still it remains difficult to decide whether to take the test or not.

A Proofs

Proof of Lemma 1. By the assumption on, U_p, we immediately obtain $u_p^* \leq 0$ and $u_p^*(0) = u_p^*(1) = 0$. Fix $a, b, \rho \in [0,1]$ and define $m = \rho a + (1-\rho)b$. Then by the optimality of y^*, we have convexity:

$$\begin{aligned}
& u_p^*(\rho a + (1-\rho)b) \\
= \quad & (\rho a + (1-\rho)b)U_p(1, y^*(m)) + (1 - (\rho a + (1-\rho)b))U_p(0, y^*(m)) \\
= \quad & \rho U_p(a, y^*(m)) + (1-\rho)U_p(b, y^*(m)) \\
\leq \quad & \rho u_p(a, y^*(a)) + (1-\rho)u_p(b, y^*(b)) \\
= \quad & \rho u_p^*(a) + (1-\rho)u_p^*(b).
\end{aligned}$$

Convexity over $[0,1]$ implies continuity over $(0,1)$. Continuity in 0 follows from $u_p^*(0) = u_p(0,0) = 0$, and from the facts that $0 \geq u_p^*(b) \geq u_p(b,0)$ and

$$\lim_{b \to 0} u_p(b,0) = \lim_{b \to 0}(1-b)U_p(0,0) + bU_p(1,0) = 0.$$

Continuity in 1 follows analogously. □

[29] Whether this is the case depends on the country. Germany, for example, does not allow health insurers to discriminate based on genetic test results. For the relevant legal guidelines, see Gendiagnostikgesetz §18.

Proof of Proposition 1. We first show that $E[V(B)] \leq \overline{V}(p)$ for all $B \in \mathcal{B}$ and then construct B^* such that it attains this upper bound. For the upper bound fix some $B \in \mathcal{B}$ and observe that since $\overline{V}$ is greater than or equal to V and concave we obtain

$$E\left[V(B)\right] \leq E\left[\overline{V}(B)\right] \leq \overline{V}(E[B]) = \overline{V}(p)$$

by Jensen's inequality. Therefore, we can at most achieve $\overline{V}$ evaluated at the prior belief p. Thus, $B^* = p$ is optimal whenever $V(p) = \overline{V}(p)$. To see that we can always achieve $\overline{V}(p)$ we construct a random variable B^* with

$$E[V(B^*)] = \overline{V}(p)$$

for the other case where $V(p) < \overline{V}(p)$. Note that by its minimality, $\overline{V}$ is linear on all open intervals J with $V(b) < \overline{V}(b)$ for all $b \in J$. Denote by $I = (b_l, b_h)$ the largest interval with the properties that $p \in I$ and $V(b) < \overline{V}(b)$ for all $b \in I$. Since this is the maximal interval, V and $\overline{V}$ must coincide in b_l and in b_h.[30] Now choose B^* as the unique random variable which takes only values b_l and b_h and which has expected value p. B^* is given explicitly in the proposition. Since V and $\overline{V}$ agree on the two values of B^* and by the linearity of $\overline{V}$ on I, we have

$$E[V(B^*)] = E[\overline{V}(B^*)] = \overline{V}(E[B^*]) = \overline{V}(p)$$

and thus B^* indeed attains the upper bound. □

Proof of Lemma 2. Applying Bayes' rule, we immediately obtain the requirements

$$P[X = 1|S = \text{“Good”}] = \frac{\alpha p}{\alpha p + \beta(1-p)} = b_h$$

and

$$P[X = 1|S = \text{“Bad”}] = \frac{(1-\alpha)p}{(1-\alpha)p + (1-\beta)(1-p)} = b_l.$$

[30] In particular, for the case of $I = (0,1)$ where this does not immediately follow from the definition of I, it is easy to check that by the minimality of $\overline{V}$, V and $\overline{V}$ *always* coincide in 0 and 1: Otherwise we could modify $\overline{V}$ on a small interval to make it smaller.

Solving for α and β yields the solution given in the proposition. It remains to check that $\alpha, \beta \in [0,1]$. For β this is clear since it is the product of two fractions which obviously lie in $[0,1]$ by $0 \leq b_l < p < b_h \leq 1$. $\alpha \geq 0$ also follows immediately. $\alpha \leq 1$ is a consequence of the fact that

$$\frac{p - b_l}{b_h - b_l} \leq \frac{p}{b_h}.$$

□

Proof of Proposition 2. Since the optimal test is invariant to multiplying V by a constant, we can reinterpret decreasing θ as adding a convex function to V. Recalling the definition of b_l and b_h as the boundaries of maximal intervals over which $\overline{V}$ strictly dominates V, the result follows from the following claim: Let f be a convex function and denote by $\overline{V+f}$ the smallest concave function greater than $V+f$. Then, if $\overline{V}$ is strictly greater than V on an open interval I, $\overline{V+f}$ is strictly greater than $V+f$ over I as well. The main step in proving the claim consists of proving the inequality

$$\overline{V}(b) + f(b) \leq \overline{V+f}(b) \tag{3}$$

for all $b \in [0,1]$. To see this inequality, fix some $q \in [0,1]$, denote by $\mathcal{B}_q$ the random variables on $[0,1]$ with mean q and denote by B_V^* a solution to $\max_{B \in \mathcal{B}_q} E[V(B)]$. Then by Proposition 1 and the convexity of f we conclude

$$\begin{aligned}
\overline{V}(q) + f(q) &= \max_{B \in \mathcal{B}_q} E[V(B)] + \min_{B \in \mathcal{B}_q} E[f(B)] \\
&\leq V(B_V^*) + f(B_V^*) \\
&\leq \max_{B \in \mathcal{B}_q} V(B) + f(B) \\
&= \overline{V+f}(q)
\end{aligned}$$

which proves (3). The claim now follows from (3) via

$$V(b) < \overline{V}(b) \quad \Rightarrow \quad V(b) + f(b) < \overline{V}(b) + f(b) \leq \overline{V+f}(b).$$

□

Proof of Proposition 3. For fixed d_l the constrained optimal test can be constructed as follows: For $d \in (p, 1]$, define g_d as the straight line connecting $(d_l, V(d_l))$ and $(d, V(d))$. For all d_h, we have $E[V(D(d_l, d_h))] = g_{d_h}(p)$ and by assumption there exists d_h such that $g_{d_h}(p) > V(p)$. Let g be the straight line through $(d_l, V(d_l))$ with the property that g has the smallest slope among all straight lines which are greater than or equal to V over $[p, 1]$. Clearly, $g(p) \geq E[V(D(d_l, d_h))]$ for all $d_h \in (p, 1]$. Moreover, by the continuity of V this inequality is an equality for some values of d_h and, accordingly, $g \equiv g_{d_h}$ for these values. Denote by d_h^* the smallest value in $[p, 1]$ such that $g_{d_h^*} \equiv g$. By assumption, $d_h^* > p$. Thus we have identified a constrained optimal belief $D(d_l, d_h^*)$ and it remains to show that $d_h^* \leq b_h$. Note first that $g_{b_h}(b) \leq g_{d_h^*}(b)$ for all $b > d_l$ by the definition of d_h^*. Denote by f the straight line connecting $(b_l, V(b_l))$ and $(b_h, V(b_h))$ and note that $f(b) = \overline{V}(b)$ for $b \in [b_l, b_h]$ and $f(b) \geq \overline{V}(b)$ for $b \geq b_h$ by the concavity of $\overline{V}$. Since $f(d_l) = \overline{V}(d_l) > V(d_l) = g_{b_h}(d_l)$ and $f(b_h) = V(b_h) = g_{b_h}(b_h)$, it follows that $g_{b_h}(b) > f(b)$ for $b > b_h$ as both functions are linear. Yet this implies that for $b > b_h$

$$g_{d_h^*}(b) > g_{b_h}(b) > f(b) \geq \overline{V}(b) \geq V(b).$$

Since $g_{d_h^*}(d_h^*) = V(d_h^*)$ we must have $d_h^* \leq b_h$. □

Proof of Example 1. We have to show that the second derivative $V''(b) = \theta U_a''(b) + (1-\theta) u_p^{*\prime\prime}(b)$ switches signs at most once and if it does then from negative to positive. Under our assumption on $U_p(x, \cdot)$, the function $u_p^*(b)$ is given by $u_p^*(b) = -b(1-b)$ and thus $u_p^{*\prime\prime}(b) = 2$ for all b. Since U_a'' is monotone, $\theta U_a''(b)$ and $(1-\theta) u_p^{*\prime\prime}(b)$ intersect at most once and it is easily checked that the resulting signs of V'' match the claims in the lemma. □

Proof of Proposition 4. The proof is organized as follows: We state a characterization of the function $\overline{V}$ in the setting of Assumption 1. This is the content of Lemma 3 below. We then show how to obtain the result the proposition from this lemma. The proof of Lemma 3 follows at the end. The construction of $\overline{V}$ is depicted in Figure 3. Define for $z \in \mathbb{R}$ the linear function $g_z : [0, 1] \to \mathbb{R}$ as the straight line connecting $(0, z)$ and $(1, V(1))$. Pick a value z^* such that g_{z^*} is tangential to V at some point $(b_t, V(b_t))$. Set $\overline{V}$ equal to V on $[0, b_t]$ and equal to g_{z^*} on $[b_t, 1]$. In the picture, it is evident that this construction yields a concave function which weakly dominates V. Lemma 3 shows that this construction always works and that the resulting function is indeed $\overline{V}$.

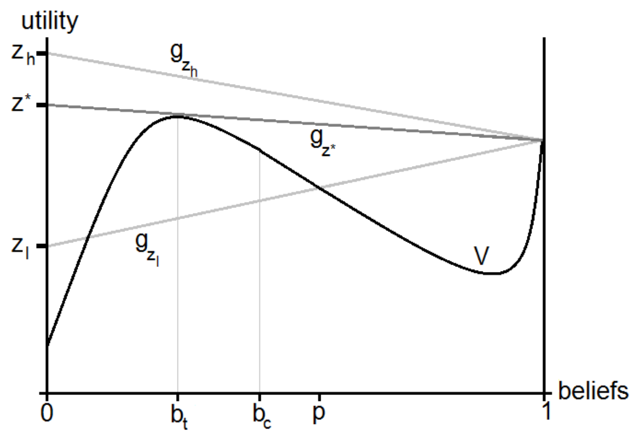

Figure 3: Construction of $\overline{V}$ under Assumption 1

Lemma 3. *Under Assumption 1, the function $\overline{V}$ can be constructed as follows.*

(i) If $g_{V(0)}(b) \geq V(b)$ for all $b \in [0,1]$, set $\overline{V} = g_{V(0)}$ and $b_t = 0$.

(ii) Otherwise, there exist a unique $z^* \in \mathbb{R}$ and $b_t \in (0, b_c]$ such that $g_{z^*}(b) \geq V(b)$ for all b and g_{z^*} is a tangent to V in b_t. Set

$$\overline{V}(b) = \begin{cases} V(b) & \text{if } b \leq b_t \\ g_{z^*}(b) & \text{if } b > b_t. \end{cases}$$

The boundary case (i) of the lemma corresponds to the situation where the straight line connecting $(0, V(0))$ and $(1, V(1))$ dominates the graph of V for all b. In this case, full revelation is optimal at all priors (and we have $b_t = 0$). Case (ii) is the one depicted in Figure 3. In that case, the function $\overline{V}$ is first concave and then linear, making the optimal test prior-dependent. For $p \leq b_t$ we have $V(p) = \overline{V}(p)$ by Proposition 3 and thus, again, the optimal test is non-revealing by Proposition 1. By Proposition 3, we also know that $\overline{V}$ is linear over $[b_t, 1]$. Thus, for $p > b_t$, a test which induces beliefs b_t or 1 attains $\overline{V}(p)$. By Proposition 1, it is thus an optimal test. The explicit probabilities α and β then follow from Lemma 2. $\square$

Proof of Lemma 3. The proof proceeds in two steps: In the first step we show that the construction of $\overline{V}$ in the statement of the proposition is always valid, i.e., that there exists a unique function which can be constructed according to the prescriptions in the proposition. Denote the constructed function by $\widehat{V}$. In the second step we verify that $\widehat{V} = \overline{V}$, i.e., that the constructed function is indeed the smallest concave function dominating V.

Step 1 : The construction of $\overline{V}$ given in the proposition is satisfied by a unique function $\widehat{V}$. Case (i) is clear so we turn to case (ii). Note that since V is a continuous function on a compact set (and thus bounded) and since its derivative in $b = 1$ must be bounded from below by strict convexity near 1, we can choose real numbers $z_l < z_h$ with the following properties: $g_{z_h}(b) > V(b)$ for all $b < 1$ and $g_{z_l}(b) < V(b)$ for some $b \in [0, 1]$. Define the compact set $Z = [z_l, z_h]$ and define z^* via

$$z^* = \inf\{z \in Z | g_z(b) > V(b)\ \forall b \in [0, 1)\}.$$

By the continuity of V and our choice of Z this infimum is actually attained. Since we are in case (ii) we also know that $z^* > V(0)$ since g_z is monotonic in z. Since g_{z^*} is defined as an infimum over all g_z which are greater than V and since g_z is continuous in z it follows that there must exist some $b_t \in (0, 1)$ for which $g_{z^*}(b_t) = V(b_t)$. Here we can exclude $b_t = 0$ since $z^* > V(0)$. g_{z^*} and V cannot cross at this intersection because otherwise we could increase z^* slightly and still have an intersection, contradicting the minimality of z^*. Thus, g_{z^*} and V must have the same slope in b_t, i.e. g_{z^*} is a tangent to V in b_t. Moreover, we must have $b_t < b_c$: Since V and g_{z^*} coincide in b_t and in 1, they must have the same average slope over the interval $[b_t, 1]$. This average slope equals their common slope in b_t where they are tangential since g_{z^*} has constant slope. This would immediately give a contradiction if we had $b_t \geq b_c$ since in that case V would be strictly convex (strictly increasing slope) over $[b_t, 1]$. The uniqueness of b_t follows from the strict concavity of V over $[0, b_c]$: A strictly concave function cannot be tangential from below to the same straight line at more than one point. Thus, we can always construct the function $\widehat{V}$ described in the proposition. The resulting function is indeed concave since it equals V on $[0, b_t] \subset [0, b_c]$ and then continues with constant slope. Moreover, by the definition of z^*, we have $\widehat{V}(b) \geq V(b)$ for all $b > b_t$.

Step 2 : $\widehat{V}$ from Step 1 is indeed the smallest concave function dominating V, $\widehat{V} \equiv \overline{V}$.

Recall that the minimum of two concave functions is again concave. Thus we must have $V(b) \leq \overline{V}(b) \leq \widehat{V}(b)$ for all $b \in [0,1]$: If the second inequality was violated at some b then $\min(\overline{V}, \widehat{V})$ would be a concave function dominating V which was strictly smaller than $\overline{V}$ at some b, contradicting the minimality of $\overline{V}$. Since V and $\widehat{V}$ coincide on $[0, b_t]$ and in 1, they must thus also coincide with $\overline{V}$ at these values. Yet on the remaining values $(b_t, 1)$, $\widehat{V}$ is linear and thus no concave function which agrees with $\widehat{V}$ at the end points $\{b_t, 1\}$ can be smaller. This proves $\overline{V}(b) = \widehat{V}(b)$ for all $b \in [0,1]$. □

Proof of Proposition 5. We have to show that

$$V''(b) = U_a''(w(b))w'(b)^2 + U_a'(b)w''(b)$$

has exactly one interior zero and is negative to its left and positive to its right. Writing $V''(b) = 0$ as

$$-\frac{U_a''(w(b))}{U_a'(w(b)} = \frac{w''(b)}{(w'(b))^2},$$

we notice that the left hand side is positive and decreasing. Furthermore, the right hand side is continuous and positive from some interior point on. It remains to show that our assumptions on $h(y) = w^{-1}(y)$ imply that $w''(b)/(w'(b))^2$ is increasing and diverges to $+\infty$ for $b \uparrow 1$. To see this, we differentiate the identity $w(h(y)) = y$ twice, to obtain $w'(h(y))h'(y) = 1$ and

$$w''(h(y))(h'(y))^2 + w'(h(y))h''(y)) = 0.$$

Using these two identities, we find that

$$\frac{w''(h(y))}{w'(h(y))^2} = -\frac{h''(y)}{h'(y)}$$

which, by strict monotonicity of h, completes the argument as $w''/(w')^2$ and $-h''/h'$ are identical up to a monotonic transformation. □

Proof of Proposition 6. The inverse h of w is given by

$$h(y) = \exp(-\lambda(-\log(y))^{\gamma})$$

where $\gamma = \frac{1}{\rho} > 1$ and $\lambda = \kappa^{-\frac{1}{\rho}}$. The condition $\kappa < \rho^{-\rho}$ becomes $\lambda\gamma > 1$. The first two derivatives of h are given by

$$h'(y) = h(y)\frac{\gamma\lambda}{y}(-\log(y))^{\gamma-1}$$

and

$$h''(y) = -h'(y)\left(\frac{\gamma-1}{y(-\log(y))} - \frac{\gamma\lambda(-\log(y))^{\gamma-1}}{y} + \frac{1}{y}\right)$$

so that the expression in brackets corresponds to $r(y) = -h''(y)/h'(y)$. As $y \uparrow 1$, the summand $(\gamma-1)/(y(-\log(y)))$ converges to $+\infty$ while the other summands in r remain bounded. Thus, $r(y) \to +\infty$ for $y \uparrow 1$ as claimed. It remains to show that r' is positive. Taking derivatives and simplifying, we see that r' can be written as

$$r'(y) = \frac{G(-\log(y))}{y\log(y)^2}$$

where

$$G(z) = (\gamma-1)(1+\gamma\lambda z^{\gamma} - z) + z \cdot (\gamma\lambda z^{\gamma} - z).$$

We thus have to show that $G(z) > 0$ for all $z \geq 0$, i.e., over the whole range of $-\log(y)$. Defining

$$g(z) = (\gamma-1)(1+z^{\gamma} - z) + z \cdot (z^{\gamma} - z) \geq 0,$$

it follows from $\gamma\lambda > 1$ that $G(z) > g(z)$ so that it suffices to show $g(z) \geq 0$. For $z \geq 1$, we have $z^{\gamma} \geq z$ which implies $g(z) \geq 0$. For the case $z \leq 1$, notice first that the function $f(z) = z^{\gamma} - z$ is convex and has its unique minimum at $z^* = \gamma^{-\frac{1}{\gamma}}$ where it takes the value

$$f(z^*) = -\gamma^{-\frac{\gamma}{\gamma-1}}(\gamma-1).$$

We thus obtain the bound

$$g(z) \geq (\gamma - 1)(1 + f(z^*)) + z \cdot f(z^*).$$

Using that $f(z^*)$ is negative, we find that this bound implies $g(z) \geq 0$ whenever

$$z \leq -(\gamma - 1)\left(1 + \frac{1}{f(z^*)}\right) = 1 + \gamma^{\frac{\gamma}{\gamma-1}} - \gamma.$$

As $\gamma > 1$ implies $\gamma^{\frac{\gamma}{\gamma-1}} \geq \gamma$, we have thus shown $g(z) \geq 0$ for $z \leq 1$ as well. $\square$

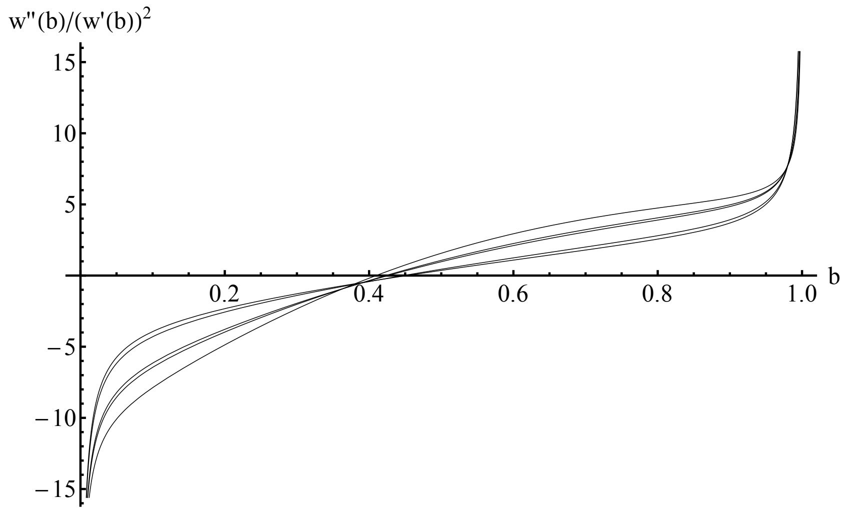

Figure 4: Graph of $w''/(w')^2$ for w of Tversky-Kahneman type, depicted for $\rho \in \{0.56, 0.6, 0.61, 0.69, 0.71\}$. The ordering corresponds to the one from top to bottom at $b = 0.8$.

References

[1] Aumann, Robert J., and Michael B. Maschler. Repeated Games with Incomplete Information. MIT Press, 1995.

[2] Baréma, Jean. The Test: Living in the Shadow of Huntington's Disease. Franklin Square Press, 2005.

[3] Baucells, Manel, and Silvia Belezza. Temporal Profiles of Instant Utility During Anticipation, Event, and Recall, Management Science, 63, 729-748, 2017.

[4] Bleichrodt, Han. Probability weighting in choice under risk: An empirical test. Journal of Risk and Uncertainty, 23, 185-198, 2001.

[5] Brunnermeier, Markus K., and Jonathan A. Parker. Optimal Expectations. American Economic Review, 95, 1092-1118, 2005.

[6] Caplin, Andrew, and Kfir Eliaz. AIDS and psychology: A mechanism-design approach. RAND Journal of Economics, 34, 631-646, 2003.

[7] Caplin, Andrew, and John Leahy. Psychological Expected Utility Theory and Anticipatory Feelings. Quarterly Journal of Economics, 116, 55-79, 2001.

[8] Caplin, Andrew, and John Leahy. The supply of information by a concerned expert. Economic Journal, 114, 487-505, 2004.

[9] Chan, Elaine, and Anirban Mukhopadhyay. When choosing makes a good thing better: Temporal variations in the valuation of hedonic consumption. Journal of Marketing Resesearch, 47, 497507, 2010.

[10] Eliaz, Kfir and Ran Spiegler. Can Anticipatory Feelings Explain Anomalous Choices of Information Sources?, Games and Economic Behavior, 56, 87-104, 2006.

[11] Epstein, Larry G. Living with Risk. Review of Economic Studies, 75, 1121-1141, 2008.

[12] Erblich, Joel, Dana H. Bovbjerg, Christina Norman, Heiddis B. Valdimarsdottir, and Guy H. Montgomery. It Won't Happen to Me: Lower Perception of Heart Disease Risk among Women with Family Histories of Breast Cancer. Preventive Medicine, 31, 714-721, 2000.

[13] Ganguly, Amanda, and Josh Tasoff. Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future. Management Science, forthcoming, 2016.

[14] Golman, Russell, and George Loewenstein, Curiosity, Information Gaps, and the Utility of Knowledge. Working Paper, 2012.

[15] Huck, Steffen, Nora Szech, and Lukas Wenner. More Effort With Less Pay: On Information Avoidance, Belief Design and Performance, Working Paper, 2015.

[16] Kadane, Joseph B., Mark Schervish, and Teddy Seidenfeld. Is Ignorance Bliss? Journal of Philosophy, 105, 5-36, 2008.

[17] Kahneman, Daniel, and Amos Tversky. Prospect theory: An analysis of decision under risk. Econometrica, 47, 263-291, 1979.

[18] Kamenica, Emir, and Matthew Gentzkow. Bayesian Persuasion. American Economic Review, 101, 2590-2615, 2011.

[19] Kemperman, Johannes H.B. The general moment problem, a geometric approach, Annals of Mathematical Statistics, 39, 93-122, 1968.

[20] Kimball, Miles S. Precautionary Saving in the Small and in the Large. Econometrica, 58, 53-73, 1990.

[21] Kőszegi, Botond. Health Anxiety and Patient Behavior. Journal of Health Economics, 22, 1073-1084, 2003.

[22] Kőszegi, Botond. Emotional Agency, Quarterly Journal of Economics. 121, 121-156, 2006.

[23] Loewenstein, George. Anticipation and the valuation of delayed consumption. Economic Journal, 97 , 666-684, 1987.

[24] Morewedge, Carey K. Utility: Anticipated, experienced, and remembered. In: Blackwell Handbook of Judgment and Decision Making, Gideon Keren and George Wu, editors, 2nd ed., Blackwell Press, 295-330, 2016.

[25] McKenna, Frank P. It won't happen to me: Unrealistic optimism or illusion of control? British Journal of Psychology, 84, 39-50, 1993.

[26] Oster, Emily, Ray Dorsey and Ira Shoulson. Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease. American Economic Review, 103, 804-830, 2013.

[27] Prelec, Drazen. The probability weighting function. Econometrica, 66, 497-527, 1998.

[28] Read, Daniel. Experienced utility: Utility theory from Jeremy Bentham to Daniel Kahneman. Thinking and Reasoning, 13, 45-61, 2007.

[29] Richter, Hans. Parameterfreie Abschätzung und Realisierung von Erwartungswerten. Blätter der DGVFM, 3, 147-162, 1957.

[30] Rosar, Frank. Test design under voluntary participation and conflicting preferences, Working Paper, 2014.

[31] Shmaya, Eran, and Leeat Yariv. Foundations for Bayesian Updating, Working Paper, 2009.

[32] Smith, James E. Generalized Chebychev Inequalities: Theory and Applications in Decision Analysis. Operations Research, 43, 807-825, 1995.

[33] Stramer, Susan L., Sally Caglioti, and D. Michael Strong. NAT of the United States and Canadian blood supply. Transfusion, 40, 1165-1168, 2000.

[34] Tversky, Amos, and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and uncertainty, 5, 297-323, 1992.

[35] Walter, Louise C., Daniel Bertenthal, Karla Lindquist, and Badrinath R. Konety. PSA Screening among Elderly Men with Limited Life Expectancies. Journal of the American Medical Association, 296, 2336-2342, 2006.

[36] Wakker, Peter P. Prospect theory: For risk and ambiguity. Cambridge University Press, 2010.

[37] Weinstein, Neil D. Why it won't happen to me: Perceptions of risk factors and susceptibility. Health Psychology, 3, 431-457, 1984.