# Optimal Revelation of Life-Changing Information*

Nikolaus Schweizer and Nora Szech

September 2013

**Abstract**

This paper studies the optimal revelation of life-changing information as in tests for severe, incurable diseases. Our model blends risk-attitudes with anticipatory utility. We characterize the optimal test design and provide conditions under which the optimal test gives either precise good news or noisy bad news, but never definite bad news. We also consider optimal test design under partial information and show how an approximately optimal dynamic unraveling of information can be implemented without any knowledge of the patient's preferences through an explicit algorithm.

**JEL Classification:** D81, D82

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

*We would like to thank Konstanze Albrecht, Doug Bernheim, Stefano DellaVigna, Péter Esö, Armin Falk, Russell Golman, Florian Herold, Rudi Kerschbamer, Emir Kamenica, Botond Kőszegi, Thomas Kruse, George Loewenstein, Matthew Rabin, Frank Rosar, 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 Mathematics, Saarland University, email: schweizer@math.uni-sb.de. Nora Szech, Department of Economics, University of Bamberg, email: nora.szech@uni-bamberg.de.

# 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. Getting more information allows one to fine-tune decisions in a better way. For example, planning the future becomes much easier once one has a concrete picture of what challenges will become relevant later on.

Imagine somebody would offer you a test that would tell you whether you would 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. Luckily, for most people, such a test seems quite hypothetical. For some people, it is however very real.

Huntington's disease is a severe hereditary genetic disorder. It starts slowly, with some coordination problems here and there. As more and more cells get damaged by the disease, 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% chance of having inherited the disease (provided that exactly one parent has it) and this is the only way the disease is transmitted. Since the 1980s a genetic test is available which allows one to perfectly determine whether a person will eventually get the disease or not.

Those who are affected often find it difficult to decide whether to take the test or not. There are books solely dedicated to this decision[1] and many "wait" for years before eventually taking the test. While the problem of testing for Huntington's disease may seem like a – disturbingly severe – minority problem, it is easy to see it as an early manifestation of a problem which will become much more wide-spread as research into human genetics progresses.

In this paper, we present a simple but fairly robust model which captures why such testing decisions are non-trivial. The model combines risk preferences,

[1] See, e.g., Baréma (2005).

for example risk aversion, with anticipatory utility. In this model, we study the structure of welfare-optimal tests which turn out to be partially but not fully revelatory in general. We then present a number of results which aim at giving guidelines for determining near-optimal tests in practice.

We propose to make use of randomization as a design tool in medical testing. This is in contrast to the usual goal in medicine to have accurate but affordable tests. The motivation comes from the substantial welfare improvements which randomized tests can yield. Thus we recommend integrating patients into the decision about how precise their test should be, a practice that would be novel to the medical field. The only flexibility medical tests offer so far is the decision whether to take them or not: The medical literature has discussed the careful use of fully revealing tests extensively[2] but has not looked into the possibility of constructing partially revealing, randomized tests. While new in medical testing, randomized mechanisms are well-established in a variety of settings, ranging from complex random procedures for determining start configurations in sports contests such as the soccer world cup to randomized pricing in the airline industry. Put differently, we emphasize in this paper that the remarkable technical progress in the possibilities for revealing information allows us to choose better amounts of information – without necessarily revealing everything.

The key idea behind our model is that the utility an agent enjoys 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 up 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 different from the classical assumption that an agent 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

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

a gamble over anticipated payoffs which – by Bayesian rationality – leaves the status quo unchanged in expectation. Thus, leaving other factors aside, an agent who is risk-averse with respect to anticipated payoffs will never want to receive any information about the future – while an agent who is risk-loving, i.e. anxious to learn about the future, would like to know everything. The risk aversion in our model is hence analogous to the standard concept of risk aversion with the only difference that it applies to anticipated outcomes instead of realized physical outcomes.

In addition to anticipatory utility our model incorporates costs: Agents with better information make better decisions. Under risk aversion regarding anticipated payoffs, this leads to a trade-off. Getting more information allows one to make better plans for the future, but it also increases the risk of obtaining bad information that will lower anticipated outcomes significantly. In this framework, which is introduced in detail in Section 2, we characterize the optimal test and show that it may be partially revealing.

In Section 3 we solve the decision problem of a doctor designing an optimal test for a patient. Our model and analysis can easily be augmented to include other behavioral factors such as curiosity or fear: To some extent, anticipatory utility merely has the role of illustrating that there are highly plausible factors against revelation of information. It is not at all necessary that it is the only such factor. Likewise, costs of making the wrong decision need not be the only factor in favor of revelation of information. We show how to design the optimal test when both types of forces stand in conflict. While similar models of anticipatory utility have previously been used to illustrate patients' ambivalent feelings about taking medical tests, our model is considerably more general than previous ones. This is necessary since the questions of test design we address are normative rather than descriptive and thus require a more robust model.

In Section 4 we develop three results which aim at giving some more concrete guidelines for practical test design. In Section 4.1 we combine decreasing risk aversion with an assumption capturing that costs of making suboptimal decisions depend on the precision of the available information rather than on its exact content. In this setting the optimal test is designed such that it sometimes delivers perfect relief but never provides extremely bad news. 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 perfect relief while you never receive information that you will die within the next $t$ years *for sure*. It seems intuitive that the decision to take this test is much easier than the decision to take a perfectly revelatory test. The same reasoning can be applied to tests for Huntington's disease.

In Section 4.2, we show how to design the best test that can be constructed based on pointwise observations of a patient's utility and cost functions. The main motivation for this analysis is that pointwise observations are the typical results of empirical methods such as questionnaires which might be used to infer what a patient wants. Section 4.3 shows how the doctor can guide the patient through a number of simple decisions until he arrives at an optimal belief. We present an algorithm which constructs this type of dynamic unraveling of information without any prior knowledge of the patient's preferences.

Section 5 provides extensions and discussions of our model: Section 5.1 provides an extension to the case where the patient's condition can take more than the two values healthy and ill. Section 5.2 addresses the time dimension of the problem. We demonstrate in a simple model that a patient's demand for information will tend to increase over time and that the optimal dynamic test will induce a gradual flow of information from the doctor to the patient. Section 6 concludes. All proofs are in the appendix.

## 1.1 Related Literature

Contributions such as Loewenstein (1987), Caplin and Leahy (2001), Brunnermeier and Parker (2005), Epstein (2008) and Golman and Loewenstein (2012) have developed concepts of anticipatory utility in the behavioral economics literature. Building on this research, Caplin and Leahy (2004), Caplin and Eliaz (2003), Kőszegi (2003, 2006) and Oster, Dorsey and Shoulson (2012) study problems of information transmission in doctor-patient relations. The main distinction between our work and the majority of these contributions is that we focus on the design of optimal tests and not on explaining features of empirically observed doctor-patient behavior.

Caplin and Leahy (2004) study testing decisions under anticipatory utility yet in the absence of instrumental information. Kőszegi (2003, 2006) blend anticipatory utility with costs of suboptimal decisions like in our framework. Kőszegi (2003) focuses on patients' preferences with regard to perfectly revelatory tests. Kőszegi (2006) studies the exchange of information between doctor and patient as a cheap-talk game where the doctor is severely limited in his power to commit on truthfulness. The patient has to choose between taking a therapy or not. For the doctor, this creates an incentive to downplay the severeness of the patent's illness: As long as the patient takes the therapy, it does not matter what the doctor tells him. As the patient understands this, the doctor can only release rough signals about the health status that are credible to the patient.[3] We abstract from such commitment problems, arguing that they can be avoided in the type of "designed communication" we are interested in, see the discussion at the end of Section 2. In an empirical study, Oster, Dorsey and Shoulson (2012) 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 the optimal design of medical tests. They focus on tests for HIV, in particular, on using partially revelatory certificates as a way to motivate agents with anticipatory utility to get tested at all. In their framework, the outcome of a matching game between certified healthy and certified possibly non-healthy people takes the role of our costs of suboptimal decisions. The authors identify a testing procedure that everybody accepts to yield an infection-free equilibrium. This design goal is different to ours as we aim at identifying the welfare-optimal test for an individual patient. Note also that establishing HIV certificates in a society is a highly complex issue from a social, ethical and technical point of view. We feel that the step from theory to practice should be much easier for the more individual-based testing problems

[3] 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.

we have in mind.

Our paper thus appears to be the first to design the welfare-optimal test in a framework where forces for and against receiving information interact in a non-trivial way. Moreover, due to the fact that we make very weak restrictions on the functional forms of preferences, we are in a position to meaningfully address issues of robustness and complexity of testing schemes: In Section 4, we discuss tests which yield good results for wide classes of preferences and show how to reach optimality without making strong assumptions on the doctor's knowledge or the patient's computation skills.

From a technical point of view, our paper is related to works on strategic conflict in information transmission: Rosar and Schulte (2012) also consider test design. The analysis of Kamenica and Gentzkow (2011) relies like ours on techniques first developed by Aumann and Maschler (1995) in the cooperative game theory literature. In fact, some of our results, notably those of Section 4, immediately yield new results in the setting of Kamenica and Gentzkow. What sets our contribution apart from these papers is that in place of problems caused by strategic interaction between economic agents we consider problems caused by the need to control one's own expectations. Put differently, the intricacies in information transmission do not come from strategic conflicts in our model but from inner conflicts in the information receiver.

# 2 The Model

Consider the following game between a receiver of information ("the patient") and a revealer of information ("the doctor") who tries to reveal what he finds out in a way that maximizes the receiver's utility. There is an initially unknown state of the world $X$ which takes the values 1 and 0 with commonly known probabilities $p$ and $1-p$. Throughout, $X = 1$ will denote the outcome preferred by the patient ("he is healthy") and $X = 0$ the other outcome ("he is ill"). The timing of the game is as follows: At the very beginning, the doctor tells the patient the distribution of a random signal $S$ which he will generate after observing $X$ and which is correlated with $X$. For instance, $S$ might be perfectly revealing, or it might be completely uninformative. Then the doctor observes the realization of $X$, generates $S$ and reveals $S$ to the patient. The patient then forms a posterior belief $B$ about the realization of $X$,

$$B = P[X = 1|S].$$

Next, the patient makes an important lifetime decision, choosing a value $y \in [0,1]$. The patient's realized utility is given by the sum of three terms: a "classical" term $U_c(X)$ for some increasing function $U_c$,[4] an anticipatory utility term $U(E[X|B])$ where $U$ is an increasing and concave function, and a cost term $-C(X,y)$ where the cost function $C : \{0,1\} \times [0,1] \to \mathbb{R}$ has the property that for fixed $x \in \{0,1\}$, $C(x,y)$ is continuous in $y$ and takes its unique minimum in $C(x,x) = 0$:

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

$\theta \in [0,1]$ is a parameter which we will later vary to investigate the interplay of the latter two terms, utility from anticipation and utility from making more informed decisions. Both, the doctor and the patient have the goal of maximizing this realized utility given the information they have.

All three terms are thought of as aggregates over all future time periods, i.e., discounted sums of future realized utilities, future anticipations and future costs of having chosen a value of $y$ which is – ex post – suboptimal. Likewise, the choice of $y$ should be understood as an aggregate over many decisions such as life insurance plans, medical insurance or occupational choice the patient must take. The key point is that the better the agent knows $X$ when he chooses $y$, the smaller is the cost term. Some more discussion of the aspect of time aggregation is found in Section 5.2 below. If we think of $X$ as something like an indicator of whether a disease will eventually break out, then the patient inevitably observes $X$ at some point in the far future. Accordingly, $y$ only captures decisions made before that point in time.

While a state variable $X$ which takes only two values is perfectly sufficient for modeling many interesting examples such as testing for Huntington's disease, we will demonstrate in Section 5.1 that our key arguments run through more or less unchanged when $X$ can take more than two values.

[4] The term $U_c$ will turn out to be spurious in the analysis below. It is included here mainly to illustrate the meaning of the two other terms.

We have formulated our model based on the story that the doctor observes the state $X$ and then uses randomization to generate the signal $S$ which is revealed to the patient. While this story has the huge advantage of being simple, it should be thought of as a placeholder for the following equivalent story which is more realistic but also more complex: The doctor himself does not observe $X$. Instead, there is a medical laboratory that observes $X$ and that has received the instructions for generating $S$ from the doctor. $S$ is generated by the laboratory which mails it to the doctor and the patient.

Taking an institution such as this medical laboratory into account is important for addressing two issues which deserve discussion: the doctor's commitment power and the effect the test result may have on the patient's preferences for information.

Concerning the commitment problem, it will obviously be hard for a doctor who knows that his patient is perfectly healthy not to tell him. However, the communication between the laboratory and the doctor can easily be anonymized in a way that eliminates problems of this kind. Kőszegi's (2006) analysis of cheap-talk in "every-day" doctor-patient relationships nicely illustrates that commitment power desirable. This affects our analysis only in so far that we assume a priori that communication is designed in a way that precludes these commitment problems.

Some authors have emphasized that preferences for information are influenced by how accessible they are and by the extent to which the decision-maker is aware of the issue at hand, see Golman and Loewenstein (2012) for a model taking these aspects into account. Indeed, a patient will likely be influenced by knowing that a sheet of paper with his diagnosis is hidden in a stack of documents right in front of him. However, it is easy to construct the testing procedure in a way that no such sheet of paper exists. The laboratory could delete all information after releasing the signal so that any further information is hidden deeply in the genetic code – just as before. Moreover, we do not think that anyone affected by Huntington's disease will ever become unaware of the issue. It is in this respect that our model *primarily* aims at "life-changing" information.

Throughout, we use the words "signal" and "test" interchangeably. The former is more in line with the language of theoretical economics and the latter is more in line with the language of the doctor-patient relationship. Thus it is useful to have both available.

# 3 Optimization

In this section we derive the optimal behavior of the doctor and the patient in our model, i.e., the SPNE of the doctor-patient game. Proposition 1 and 2 characterize, respectively, the beliefs induced by an optimal test and the optimal test itself. Proposition 3 characterizes how the optimal test becomes more revelatory if the importance of the costs of making wrong decisions increases relative to the importance of anticipatory utility.

We begin with the second decision, 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 of cost minimization is given by

$$\min_y c(b, y) \text{ where } c(b, y) = b\, C(1, y) + (1 - b)\, C(0, y).$$

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

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

Our first result, which is proved in the appendix, shows that $c^*$ is concave in $b$. This reflects the fact that more diffuse beliefs make it more difficult to reduce costs.

**Lemma 1.** *The function $c^*(b)$ is continuous and concave in $b \in [0, 1]$.*

We now turn to the doctor's problem of designing the optimal test. We take a somewhat indirect approach here by determining first the optimal belief $B^*$ induced in the patient and then constructing a test that induces this belief. To this end, denote by $\mathcal{B}$ the set of random variables valued in $[0, 1]$ which have mean $p$. By Bayesian consistency, it is clear that the doctor cannot induce any belief $B$ which is not an element of $\mathcal{B}$, i.e., which is inconsistent with the prior $p$.[5]

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

We can state the doctor's problem of maximizing the patient's expected utility conditional on the patient taking the optimal decision $y^*$ after the test as

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

Here, we have ignored the term $E[U_c(X)]$ since it is independent of the decision and we have used that $E[X|B] = B$ and thus $U(E[X|B]) = U(B)$. By assumption, $U$ is concave. Moreover, we have just seen that $-c^*$ is convex. Thus, for $\theta \in (0,1)$ the function $V$ is generally continuous but neither convex nor concave. This "conflict" between the anticipatory utility term $E[U(B)]$ and the more standard "preference for early resolution of uncertainty" term $E[-c^*(B)]$ lies at the heart of our model.

Note that concavity of $U$, i.e. risk aversion with respect to anticipated outcomes, is not necessary for this conflict to arise at least for some priors: It is sufficient that $U$ is not convex. Related to this is a big advantage of the fact that the analysis which follows now makes hardly any restrictions on $V$: All the results are still valid if we add further psychological factors such as anxiety, curiosity, fear etc. to the patient's objective function. The sole property of $V$ that is used in the following is that it is a continuous function.[6] For the sake of concreteness, we could add a term $\gamma F(b)$ modeling curiosity to 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 $\gamma$ the incentives for receiving as much information as possible become dominant).

A similar class of optimization problems was recently studied by Kamenica and Gentzkow (2011) in the context of *strategic* conflicts in information transmission. Therefore, we prefer a short and non-technical exposition of how to solve (1) which is given in the proof of Proposition 1. We refer the reader to Kamenica and Gentzkow (2011) for a more detailed presentation of this result. These techniques also have a history in cooperative game theory, see Aumann and Maschler (1995).

[6]This continuity is a highly convenient assumption since it implies that $V$ attains intermediate values, maxima and minima. It can however be relaxed at the expense of somewhat more complicated statements of the results. Notably, our behaviorally intricate results do not rely on any "irregularities" in $V$ such as kinks.

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 weakly greater than $V$. Moreover, the optimal test can be read off from the graph of $\overline{V}$ as is depicted in Figure 1: 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 $(d_h, V(d_h))$ in the figure 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 $V$, it is fairly intuitive that $\overline{V}(p)$ is exactly what the optimal test can achieve. The proof of Proposition 1 makes this point in more detail and shows in particular 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

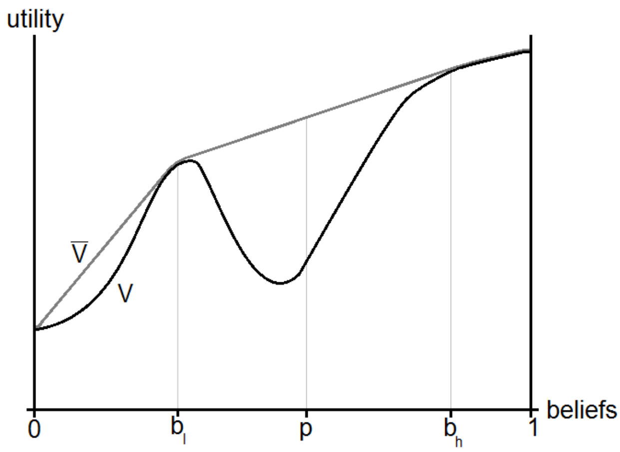


Figure 1: Construction of $\overline{V}$

$1 - p$ and $p$. Thus, in this case

$$B^* = P[X = 1|S^*] \in \{0, 1\}$$

reflects the beliefs of a perfectly informed patient who always knows whether $X = 0$ or $X = 1$ conditional on his (optimal) signal $S^*$. 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 $p$ so that the optimal signal must be uninformative. Note also that in the case where $V$ and $\overline{V}$ coincide only on some interval, it depends on the value of $p$ whether $B^*$ is informative or not.

The proposition gives a fairly general characterization of the optimal belief system. An important consequence of the proposition is that there is always a solution $B^*$ to (1) which lies in the set $\mathcal{B}_2 \subset \mathcal{B}$ 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 complete our analysis of the doctor's problem it suffices to show that for any $B \in \mathcal{B}_2$ there exists a signal which induces it. This is the result of the following proposition:

**Proposition 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*

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

*Then* $B = P[X = 1|S]$ *only takes values in* $\{b_l, b_h\}$ *and* $E[B] = p$ *where the expectation is taken over* $S$*.*

Here, $S =$ “Good” is better news than $S =$ “Bad” since it induces the higher posterior probability $b_h$ of $X = 1$. It is straightforward to rewrite the test of Proposition 2 in a way that $X$ only needs to be observed with some probability. Notably, this probability becomes small if $\alpha \approx \beta$ so that the test has little predictive power. Such a formulation becomes attractive if we take into account the costs of observing $X$ since it induces a correlation between the costs of a test and its predictive power. We assume these costs to be negligible in comparison and thus stay with the analytically convenient formulation of Proposition 2 in the following.

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

**Proposition 3.** *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,* $\theta$ *and* $\theta'$*. Then* $b_l \geq b_l'$ *and* $b_h \leq b_h'$*. Thus the optimal test under* $\theta'$ *leads to beliefs which are closer to knowledge of* $X$ *than the optimal test under* $\theta$*.*

Finally, we give a result which will become important later in Section 4.3 but which also has some intrinsic interest since it shows that the patient's preferences over tests have more structure than one might expect given that $V$

is an arbitrary continuous function: 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 $d_h^*$? We show 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 4.** *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 $D(d_h)$ the random variable with mean $p$ which takes only values $d_l$ and $d_h$. Assume there exists $d_h$ such that $E[V(D(d_h))] > E[V(p)]$ so that some beliefs $D(d_h)$ are better than no information. Then, if $d_h^*$ is a solution to*

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

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

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

# 4 Designing Good Tests

In practice it may be difficult to observe the function $V$ in its entirety – even for the patients. Moreover, given that $V$ is known, constructing the optimal test is a fairly complex two-dimensional optimization problem that will be quite a challenge to most patients. This section contains three rather different approaches to solving this problem. In Section 4.1, we show that the structure of optimal tests becomes even simpler when we make some more restrictions on the functions $U(\cdot)$ and $C(\cdot, \cdot)$, reducing the design of optimal tests to determining a single parameter. This investigation aims at developing two or three tests which have the potential of being convincing alternatives to a perfectly informative or perfectly uninformative test.[7] In Section 4.2, we derive the best test that can be constructed using only some point-wise observations of the function $V$. This result gives a partial answer to the question of how to approximate the smallest concave function greater than $V$ in practice. In Section

[7] In practice, the uninformative test, of course, takes the form of the patient deciding against taking any test.

4.3 we show how to gradually guide the patient to his optimal belief through an unraveling procedure that can be constructed without any knowledge of the patient's preferences.

### 4.1 Accuracy On Good News

The aim of this section is to show that under natural assumptions, the optimal test takes a particularly simple form: The test sometimes perfectly reveals $X = 1$ but it never perfectly reveals $X = 0$. In the language of Proposition 2 the optimal test is characterized by $\alpha \in (0, 1)$ and $\beta = 0$. In the language of the disease, if the patient receives a "Good" from the test then he knows for certain that he is healthy. If he receives a "Bad" he must correct his probability of being healthy downwards but – unlike with a perfectly revealing test – not to zero.[8]

**Assumption 1.** *$U$ is twice continuously differentiable with $U''$ being a strictly increasing function. The functions $C(x, y)$ are of the form*

$$C(x, y) = (x - y)^2$$

*for $x \in \{0, 1\}$.*

The basic idea behind the assumption is that $U$ is governed by some notion of decreasing risk aversion[9] while the costs of making the wrong decision tend to be more symmetric around what would have been right. Thus we rather expect to find negative values of $V''$ at small values of $b$ than at larger ones.

In order to construct these optimal tests we have to determine the function $\overline{V}$. We begin with an observation about the form of $V$:

[8]A similar class of tests was found optimal in Rosar and Schulte (2012) 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.

[9]Recall that $U$ is concave and thus $U''$ being increasing means that $U''(b)$ is closer to zero for larger $b$. The assumption of an increasing second derivative of $U$ was coined "prudence" by Kimball (1990). It is a necessary condition for weakly decreasing absolute risk aversion and thus satisfied by many of the standard utility functions. We share this assumption with the literature on intertemporal transfer of utility initiated by Leland (1968).

**Lemma 2.** *Under Assumption 1, 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* $\theta \in (\theta_l, \theta_h)$, *there exists a point* $b_c(\theta) \in (0, 1)$ *such that* $V(b)$ *is strictly concave on* $[0, b_c(\theta)]$ *and strictly convex on* $[b_c(\theta), 1]$. *Moreover,* $b_c(\theta)$ *is increasing in* $\theta$.

Intuitively, the lemma states that for intermediate values of $\theta$ the anticipatory utility term is dominant for small $b$ and the cost term is dominant for larger $b$. The reason for this lies in our assumption that risk aversion is more pronounced at pessimistic beliefs.

Let us briefly touch upon the robustness of the argument behind the lemma: Clearly, the conclusion of the lemma is still valid at least for most values of $\theta$ if $c^{*\prime\prime}$ does not vary too strongly. Moreover, this robustness is greater if the monotonicity of $U''$ is more pronounced, i.e. if risk aversion decreases more strongly.

In the following, we will ignore the cases where either of the two terms dominates (implying convexity or concavity of $V$) since this leads to perfectly revealing or non-revealing optimal tests as discussed in the previous section. Consequently, for the remainder of Section 4.1 we replace Assumption 1 by the following assumption which excludes these trivial cases and is less restrictive otherwise:

**Assumption 2.** *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* $[b_c, 1]$.

We construct the function $\overline{V}$ in two steps: We first show how to construct a candidate function $\widehat{V}$ which equals $V$ up to some point and then continues linearly. Then we show that this function $\widehat{V}$ equals $\overline{V}$.

The construction of $\widehat{V}$ is depicted in Figure 2. 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 $\widehat{V}$ 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 $V$. The next proposition shows that this construction always works.

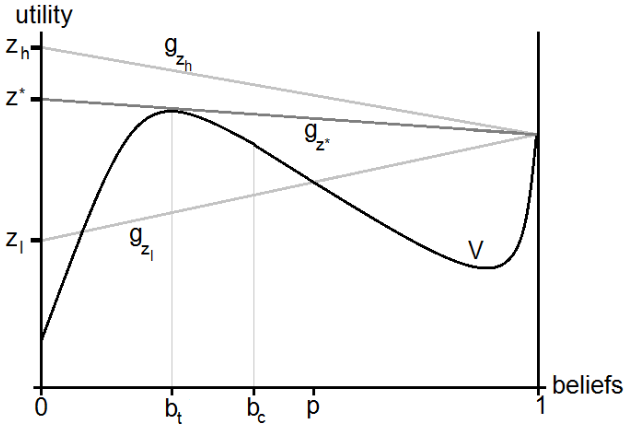


Figure 2: Construction of $\overline{V}$ under Assumption 2

**Proposition 5.** *A concave function $\widehat{V}$ which weakly dominates $V$ can be constructed as follows:*

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

*(ii) Otherwise, there exist 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*

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

The next step shows that this function $\widehat{V}$ is indeed the smallest concave function which dominates $V$.

**Proposition 6.** *We have $\overline{V} \equiv \widehat{V}$, i.e., $\widehat{V}$ is the smallest concave function with $\widehat{V}(b) \geq V(b)$ for all $b \in [0,1]$.*

Combining the preceding analysis with the result of Proposition 2, we see how to design optimal tests under Assumption 2:

**Corollary 1.** *In the notation of Proposition 2, the optimal signal is given as follows:*

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

*(ii) If $p > b_t$, then the optimal belief $B^*$ takes only values $b_l = b_t$ and $b_h = 1$. Thus, the optimal signal sometimes reveals $X = 1$ but never $X = 0$. Precisely,*

$$\alpha = \frac{1}{p}\frac{p - b_t}{1 - b_t}$$

*and $\beta = 0$.*

Recall that in proceeding from Assumption 1 to Assumption 2 we excluded some trivial cases where the optimal test is perfectly revealing – yet in most models of interest those tests will also arise as the optimum for some parameter constellations.

We end this section with the following policy implication: We have seen that in a somewhat restrictive but natural class of models the optimal signals are characterized by $\alpha \in [0,1]$ and $\beta = 0$. Thus, if a doctor wishes to include some promising tests in his portfolio of options offered to the patient – augmenting the revelatory and non-revelatory tests represented by $\alpha \in \{0,1\}$ and $\beta = 0$ – it is a good starting point to include the following three tests: $\alpha \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$ and $\beta = 0$. This even partition of $[0,1]$ can be adapted if a particular one of the five values of $\alpha$ turns out to be most popular.[10]

## 4.2 Test Design With Partial Observations

The analysis of this section takes a different approach to the problem of the doctor not knowing the function $V$ precisely. We assume that the doctor has observed $V$ on a finite subset $G$ of $[0,1]$. Our main result gives the optimal test under this limited information. There are three motivations for considering this problem: For one thing, point-wise observations are the most typical results of empirical investigation methods such as questionnaires. For another, we saw in Section 3 that even the optimal test is based on only two values of $V$ directly. Thus we have reason to hope that tests that are based on point-wise observations of $V$ may perform rather well. Indeed, we will see in a moment that observing $V$ at only four points is sufficient for constructing a test which is provably better than both, perfect revelation and no revelation at all. Finally,

[10]One can also argue that higher values of $\alpha$ are more reasonable than lower ones: If we set $p = \frac{1}{2}$ as in Huntington's disease and set $b_t$ to the middle of $[0,p]$, i.e. $b_t = \frac{1}{4}$ then we obtain $\alpha = \frac{2}{3}$. Thus $\alpha \in \{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}\}$ might be the more promising starting point.

we rely on this section when constructing the dynamic testing procedure of Section 4.3.

Before we state our central result, determining the restricted optimal test, let us define our optimality criterion: By "the optimal signal which can be constructed from knowing $V$ only on the set $G$" we mean that the test maximizes the patient's expected utility among all signals for which the patient's expected payoff can be calculated using only evaluations of $V$ on $G$.

**Proposition 7.** *Let $G \subset [0,1]$ be a finite set such that $\min G < p < \max G$. Let $V_G$ be the smallest concave function over $[\min G, \max G]$ such that for all $b \in G$ we have $V(b) \leq V_G(b)$. Then a signal attaining $V_G(p)$ can be constructed as follows: Let $(b_l, b_h) \subset [\min G, \max G]$ be the largest interval containing $p$ over which $V_G$ is linear and implement the signal from Proposition 2 with these values of $b_l$, $b_h$ and $p$. This is the optimal signal which can be constructed from knowing $V$ only on the set $G$. If in addition $\{0, p, 1\} \subset G$ then $V_G(p) \geq \max(E[V(X)], V(p))$. Thus, in this case the signal is weakly better than both perfect revelation and no revelation.*

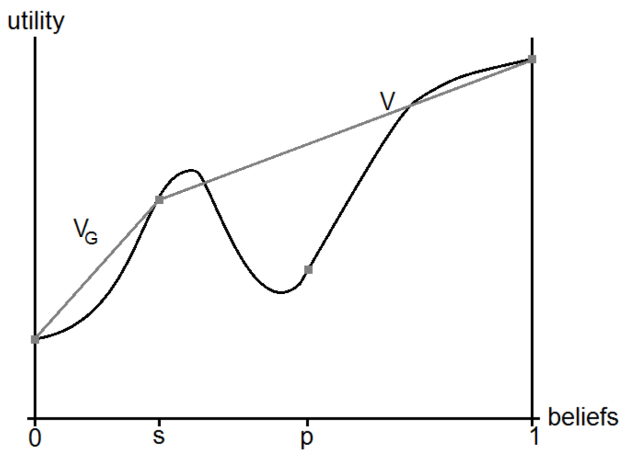


Figure 3: Construction of $V_G$ for $G = \{0, s, p, 1\}$. The optimum is characterized by $b_l = s$ and $b_h = 1$ in this case.

The construction of $V_G$ is depicted in Figure 3. Note that $V_G$ is the upper contour of the convex hull of the points $(b, V(b))_{b \in G}$. Since computing convex hulls of finite sets is an extremely well-studied problem in applied mathematics, Proposition 7 can be seen as first step towards the resolution of a problem pointed out by Kamenica and Gentzkow: Especially if $X$ takes more than two values (see Section 5.1), computing the function $\overline{V}$ is non-trivial even if $V$ is known. Proposition 7 shows that for discrete approximations of $V$, this problem is reduced to a well-studied problem. Imposing some regularity on $V$, we conjecture that one can get a handle on the approximation error $|V_G(p) - \overline{V}(p)|$ as the number of points in $G$ gets large, but we will not follow this direction further here.

## 4.3 Dynamic Unraveling Of Information

The previous two sections studied how much a doctor who has only limited information can achieve in test design. One main motivation for these considerations was the observation that designing the optimal test himself may be prohibitively difficult for a patient. In this section we address the same problem from another perspective: Can we guide the patient through a number of simple choices until he arrives at an (approximately) optimal test?

The sequence of tests we propose has the following properties: The patient is sequentially offered one test at a time and decides whether he wants to take this test before the next offer or not. Information about which tests the patient has previously accepted and taken is used in the design of the next test. This helps to keep the testing procedure short. All the doctor needs to know to design the sequence of tests is the prior belief $p$.

The idea that such a testing procedure is desirable is driven by a notion of complexity of mechanisms which we keep informal but which we think is convincing: It is good to offer only few alternatives simultaneously, it is good to keep the overall testing procedure short, and it is better to offer real choices than hypothetical ones. As will become clearer below, we face a trade-off between an approximation error and complexity: A smaller error leads to a longer testing procedure and more narrow – and thus difficult – decisions.

The key difficulty in designing such a dynamic testing procedure is that we must be careful that the patient does not learn too much at any point in time, i.e., that the patient's posterior belief never leaves the interval $[b_l, b_h]$ of the optimal static test. If the doctor gives an additional piece of information to a patient who already knows too much there is some chance that it will confuse him but there is *always* the chance that he learns even more.

How can a doctor make sure that he never tells the patient too much without knowing the values of $b_l$ and $b_h$? The key idea is to increase the amount of information gradually – yet this is not trivial since $b_l$ and $b_h$ need not lie symmetrically around $p$. Thus a test which is too informative in one direction can be less informative than optimal in the other. The observation that makes designing our dynamic test possible is the following:

**Corollary 2.** *Fix some $d_l \in (b_l, p)$. Using the notation of Proposition 4, we then have*

$$E[V(D(d_h))] \leq E[V(D(b_h))]$$

*for all $d_h \geq b_h$.*

The corollary states that a test which induces beliefs in $\{d_l, b_h\}$ where $d_l$ is less informative than optimal and $b_h$ is optimal, is preferred by the patient to any test which induces the same $d_l$ and a value $d_h \geq b_h$ that is more informative than optimal. As we will see in the following, this result implies, both, that we can avoid going through all the available tests, and that the patient will not accept a test which is too informative if tests are offered to him in the right order.

The framework of this section is similar to the one of Section 4.2 but formulated in a way that the results of Section 3 can be applied directly. Fix some integer $n > 2$, let $\varepsilon = \frac{1}{n}$ and let $G = \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$ be a discretization of $[0, 1]$. Assume that the function $V$ is continuous, piece-wise linear, and linear between any two adjacent values in $G$. $V$ should be thought of as a discrete approximation of some "true" $\widetilde{V}$ which is interpolated linearly between the points in $G$. Reformulating the model in terms of this discretized $V$ allows us to do the following: We can determine the approximate optimum in the sense of Section 4.2 simply by applying the "perfect" optimization of Section 3 to this modified $V$. In short, we offer a test which is optimal if the patient's preferences are linear between the values of $G$. Thus we typically make a small discretization error. Yet since we will only offer tests which induce beliefs with values in $G$, this possible discrepancy between our $V$ and the true $\widetilde{V}$ will not influence the patient's decisions in the testing procedure.

Under our assumption on $V$, it follows like in Section 4.2 that the two values $b_l$ and $b_h$ of an optimal test must lie in $G$. For convenience, we also assume that the prior $p$ lies in $G$.[11] We are now ready to formulate our dynamic testing procedure. We use two sets of values $\{d_l, d_h\}$ and $\{s_l, s_h\}$ in the test: The former values contain the beliefs induced by the test currently offered to the patient. The latter values are used for saving an interval $(s_l, s_h)$ which can be excluded from the further search for the optimal test.

**Proposition 8** (Algorithm)**.** *Conduct the following steps until the patient holds a belief in $\{0, 1\}$ or until the patient has been offered the perfectly revealing test and refused it:*

*(i)* *Set $s_l = d_l = p - \varepsilon$ and $s_h = d_h = p + \varepsilon$.*

*(ii)* *Offer the patient a test which induces beliefs only in $\{d_l, d_h\}$. Conduct the test if the patient accepts the offer.*

*(iii)* *If the patient refused the last offered test, increase $d_h$ by $\varepsilon$. If this implies $d_h > 1$, set $d_h = s_h$ and decrease $d_l$ by $\varepsilon$. Go back to Step (ii).*

*(iv.1)* *If the patient accepted the last offered test and his posterior belief is $d_h$, set $p = d_h$, increase $d_h$ by $\varepsilon$, set $s_h = d_h$ and $s_l = d_l$ and go to Step (ii).*

*(iv.2)* *If the patient accepted the last offered test and his posterior belief is $d_l$, set $p = d_l$, decrease $d_l$ by $\varepsilon$, set $s_l = d_l$ and $s_h = d_h$ and go to Step (ii).*

*At the end of this testing procedure, the patient either holds belief $b_l$ or $b_h$. Thus, the procedure is equivalent to offering the patient the optimal test.*

Note that the testing procedure is constructed in a way that if the patient accepted a test with parameters $\{d_l, d_h\}$, the interval $(d_l, d_h)$ is eliminated from the further search for an optimal test. This substantially decreases the number of tests which have to be offered to the patient. Of course, the testing procedure may still take rather long for small $\varepsilon$. Thus it is important to note that the procedure can be stopped without harm at any point – and possibly be continued later on. For actually implementing this type of testing procedure it seems most promising to set up a computer program that contains the realization of $X$ and makes the subsequent offers to the patient.

[11] If $p$ does not lie in $G$, we can simply add $p$ to $G$ and set the two values of the initial test to the elements in $G$ adjacent to $p$ in the testing procedure of Proposition 8.

One could try to avoid the discretization error by constructing a continuous-time testing procedure: Once we offer the patient a discrete sequence of choices we can never hope to exactly determine a continuous piece of information such as the structure of the optimal test. Yet we feel that there is little reason to believe that choosing $\varepsilon$ smaller than 0.01 and thus varying probabilities in steps of less than one percent leads to any improvements which are still perceived by patients. Finally, the procedure can immediately be generalized to sets $G$ where points are not at an equal distance. For instance we could choose finer discretizations near 0 and 1.

# 5 Extensions

## 5.1 Many States Of The World

In this section we demonstrate that most of our results easily generalize to more general random variables $X$.[12] To begin with, let us however emphasize once more that the case of $X \in \{0,1\}$ is perfectly sufficient for many applications. Assume now that $X$ takes values in a finite set $\mathcal{X} = \{x_1, \ldots, x_n\} \subset [0,1]^d$ with probabilities $p = (p_1, \ldots, p_n) \in \Delta_n$ where $\Delta_n$ denotes the probability distributions on $\mathcal{X}$.

There is only one major adjustment we have to make to generalize our analysis to this model: In the case where $\mathcal{X} = \{0,1\} \subset [0,1]$ we had a one to one correspondence between beliefs and conditional expectations of $X$ through $E[X|B] = B$. In the general case, many posterior beliefs lead to the same conditional expectation of $X$ and thus we have to work with the posterior beliefs more directly. To this end, define $\mathcal{B}$ as the set of random variables on $\Delta_n$ with mean $p$ and let $B \in \mathcal{B}$ be the belief induced in the agent by the signal

$$B = (B_1, \ldots, B_n) \in \mathcal{B} \text{ where } B_i = P[X = x_i|S].$$

In this framework, we can solve the maximization problem

$$\max_{B \in \mathcal{B}} E[V(B)] \tag{2}$$

[12] Exceptions are discussed at the end of this section.

for any $V : \Delta_n \to \mathbb{R}$ going through exactly the same steps as in Proposition 1. The only difference is that now $\overline{V}$ is the smallest concave function dominating $V$ on $\Delta_n$ and not on $[0, 1]$, see Kamenica and Gentzkow (2011) for details. Again, $\overline{V}$ is linear wherever it strictly dominates $V$ and thus an optimal test can in principle be calculated as in Proposition 2. As Kamenica and Gentzkow (2011) point out, getting a handle on $\overline{V}$ is generally non-trivial if $X$ takes more than two values. We thus propose to choose $G$ as a suitable discretization of $\Delta_n$ and calculate a discrete approximation to the optimal test as outlined in Proposition 7. This approach has the added advantage that the structure of $V_G$ is comparatively simpler than that of $\overline{V}$ since the former is derived from the convex hull of a finite set of points. This considerably simplifies deriving the restricted optimal test from $V_G$.

We still have to justify that (2) is the right generalization of the doctor's optimization problem. If we define $c^*(b) : \Delta_n \to \mathbb{R}$ analogously to the case of $X \in \{0, 1\}$ we still obtain a concave function by an argument entirely parallel to Lemma 1. To see that a concave function $U : \Delta_n \to \mathbb{R}$ is in line with our previous assumptions, we observe the following:

**Lemma 3.** *Let $u : [0, 1]^d \to \mathbb{R}$ be a concave function. Then $U : \Delta_n \to \mathbb{R}$ defined by*

$$U(b) = u(E[X|B = b])$$

*is concave as well.*

Thus we see that concavity in conditional expectations implies concavity in posterior beliefs. Accordingly, the conflict between concavity of anticipatory utility and concavity of costs is present in this more general model as well.

How to generalize Section 4.1 is less obvious and we will not dig deeper here – but, of course, the basic principle that the optimal signal is less revealing in regimes where risk aversion is more pronounced should be universal. Thus under decreasing risk aversion, tests which are more revelatory with respect to good news than with respect to bad news are still promising candidates for good tests. A similar remark applies to Section 4.3.

## 5.2 Time

In this section, we address the time dimension of the problem more explicitly. We first argue that a patient's demand for information should be increasing over time in a simple framework. Then we turn to an elementary model where we allow for multiple tests at different time points. We close with some more general remarks on dynamic welfare maximization under anticipatory utility.

Our argument for an increasing demand for information is based on the following observations: In the problems we are interested in, there is typically a time of anticipation followed by a period of bearing the costs of having made wrong decisions. The terms $U$ and $c^*$ in our model thus belong to different time intervals where the time of $U$ gets shorter over time while the time of $c^*$ comes nearer. Accordingly, we should expect the weight of the cost term to increase over time since the costs are discounted less heavily, and we should expect the weight of the anticipation term to decrease since the number of time periods where anticipation happens decreases. To formalize this in a simple model inspired by Loewenstein (1987), consider discrete time periods $s = 0, 1, \ldots$, let $0 < \delta < 1$ be a discount factor and denote by $T$ the time when the disease breaks out.[13] Then we define the utility function $V_s$ for time $s = 0, \ldots, T$ as

$$V_s(b) = \sum_{k=0}^{T-s} \delta^k U(b) - \sum_{k=T-s+1}^{\infty} \delta^k c^*(b).$$

Normalizing $V_s$ in terms of weights $\theta_s$ and $1 - \theta_s$, it is easy to see that $\theta_s$ is decreasing in $s$. Accordingly, the cost term will become more important over time, and taking a more informative test becomes more attractive as time goes by, see Proposition 3. Of course, this is a very simple model – yet it contains two natural arguments for an increasing demand for information. Thus, we expect this monotonicity behavior to be fairly robust also in more general models.

Now consider a model where a patient and a doctor meet every period $s = 0, \ldots T - 1$ for a test. The doctor's objective function for designing the at time $s$ is given by $V_s(b) = \theta_s U(b) - (1 - \theta_s)c^*(b)$ for a sequence of coefficients $\theta_s$ and concave functions $U$ and $c^*$. Thus for all $s$, the doctor optimizes the

[13] Allowing $T$ to be random would not change the argument.

patient's utility as perceived now and neither the doctor nor the patient take into account the effect of future tests on the patient's beliefs. Denote by $p(s)$ the patient's belief before the test at time $s$. The numbers $p(s)$ are common knowledge, e.g., the patient reports them to the doctor. $p(0) = p_0$ is a constant. Denote by $b_l(s)$ and $b_h(s)$ the values of an optimal belief for a patient with prior $p(s)$, calculated from the smallest concave function $\overline{V}_s$ dominating $V_s$ as described in Section 3.

Note first the following stability property of the optimal tests from Section 3: If the sequence $\theta_s$ is constant, then a patient who received the optimal test at some time $s$ does not desire any further tests at later times. The reason for this is that for both values of an optimal belief for a patient with prior $p$, $b_l$ and $b_h$, we must have $V_s(b_l) = \overline{V}_s(b_l)$ and $V_s(b_h) = \overline{V}_s(b_h)$ by the construction of $b_l$ and $b_h$. Thus the optimal test for a patient whose prior beliefs are concentrated on these two values is perfectly non-revealing. This is the case for a patient who once took the optimal test.

Now assume that the sequence $\theta_s$ is decreasing as in the previous model. Then we observe that the patient's sequence of a beliefs has a particularly simple structure: The patient's belief at a fixed time $s$ can only have one of two values regardless of the sequence of test outcomes up to that point:

**Proposition 9.** *The values of $b_l(s)$ and $b_h(s)$ are independent of previous test results and correspond to the boundaries of the largest open interval around $p_0$ on which $V_s(b) < \overline{V}_s(b)$. The prior $p(s)$ has the value $b_l(s-1)$ if the test at time $s-1$ provided bad news and the value $b_h(s-1)$ if that test provided good news.*

Consequently, there are only two different optimal tests at time $s$ depending on whether the patient has the pessimistic prior $p(s) = b_l(s-1)$ or the optimistic $p(s) = b_h(s-1)$. By Proposition 3, these priors gradually move closer to 0 and 1 as time passes, reflecting a gradual transmission of information. As the end time $T$ comes near we will typically have $b_l(s) \approx b_l(s-1) \approx 0$ and $b_h(s) \approx b_h(s-1) \approx 1$ so that the optimal test is nearly revealing. However, since the patients prior is also close to perfect knowledge of $X$, the amount of information exchanged through the optimal test is tiny with high probability – only with a very small probability a pessimistic patient becomes optimistic and vice-versa.[14]

All these considerations are not meant to obscure the fact that a fully-fledged dynamic analysis of a model including the effect of future tests, a dynamic choice of $y$ and higher order anticipatory utility[15] is a challenging topic for future research.

Besides the usual technical problems of handling dynamic models, a central difficulty in our setting is that anticipatory utility and related concepts tend to induce time-inconsistencies in agents, see, e.g., Kőszegi (2010). We are convinced that such time-inconsistencies should not be viewed as weaknesses of the model but rather as stylized facts of human nature which have been observed in a multitude of ways. Nevertheless, it is a delicate question – far beyond the scope of the present paper – how a benevolent doctor should maximize the utility of a time-inconsistent patient. To this end, let us remark three things:

(i) From a descriptive point of view, a doctor who wishes to maximize a patient's utility as it is perceived by the patient right now seems like a fairly convincing description of the objectives of real-life doctors – anything else seems too complicated to be of practical importance. This is especially true for specialized doctors who do not have regular contact with a patient before a particular disease actually breaks out. The latter case captures many of the examples we have in mind, including testing for Huntington's disease.

(ii) From a normative point of view, the time-aggregate utility function $V$ we considered above could be viewed as both – taking into account future changes in the patient's utility or not.[16] In addition, a patient's future preferences will be fiendishly hard to observe. Thus postulating that welfare maximization should take into account future changes in

[14]These last claims are easily confirmed by inspecting Proposition 2.

[15]By higher order anticipatory utility we mean utility from "looking forward to looking forward to looking forward to ..." which is ignored in our analysis just like first order utility from "looking forward to ..." is ignored in conventional models. At this point however, we do not expect strong qualitative changes in the predictions when moving to higher orders.

[16]In the former case, however, the patient's and doctor's interest will no longer be perfectly aligned and a strategic component would enter optimal test design. As Kamenica and Gentzkow (2011) have shown, the model framework we have here can handle strategic interaction as well, see also Caplin and Leahy (2004).

preferences would mean the end to many attempts at dynamic welfare maximization.

(iii) As we have argued above, the patient's demand for information will tend to increase over time. Thus giving the patient as much information as he desires at each point in time is a promising proxy for an optimal dynamic revelation strategy regardless of the precise optimality criterion.

# 6 Conclusion

One might argue that there are moral obligations which force doctors to give patients as much information as possible. Yet one could argue likewise that a doctor or an institution that has the power to design a welfare-optimal test is under a moral obligation to do so. Even if the exact specification of the optimal test is difficult, patients can only gain from choosing their test out of a menu. Put differently, offering nothing but the perfectly revealing test is only justified if it is clear that people want to know the good and the bad for sure. Yet this is refuted by the difficulties patients face when deciding to take the test for Huntington's disease.

We have provided a simple model combining anticipatory utility and risk preferences as a basis to work on test design for life-changing information. Anticipatory utility is known to play a significant role when it comes to important events, both, on shorter and longer time horizons. As risk preference, we assume risk aversion, since we feel this is the right assumption when physical decay and death are possible outcomes. Yet our model provides the flexibility to combine anticipation with any other kind of utility from future events. We could for instance add a component of risk-loving behavior, hence curiosity about the future. Our model shows that incentives for acquiring and not acquiring information do not simply cancel out, and accordingly, it is not simply such that the stronger of the two effects wins. This gives rise to optimal tests which are partially revealing.

With regard to Huntington's disease, some may think about more conventional economic explanations for not taking the test. For instance, one might invoke the costs of the test or the difficulties of finding a health insurer if the test result is bad.[17] 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 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.* Fix $a, b, \rho \in [0,1]$ and define $m = \rho a + (1-\rho)b$. Then by the optimality of $y^*$, we have concavity:

$$\begin{aligned} & c^*(\rho a + (1-\rho)b) \\ = \; & (\rho a + (1-\rho)b)C(1, y^*(m)) + (1 - (\rho a + (1-\rho)b))C(0, y^*(m)) \\ = \; & \rho c(a, y^*(m)) + (1-\rho)c(b, y^*(m)) \\ \geq \; & \rho c(a, y^*(a)) + (1-\rho)c(b, y^*(b)) \\ = \; & \rho c^*(a) + (1-\rho)c^*(b). \end{aligned}$$

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

$$\lim_{b\to 0} c(b,0) = \lim_{b\to 0}(1-b)C(0,0) + bC(1,0) = 0.$$

Continuity in 1 follows analogously. □

*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 weakly greater than $V$ 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

[17] Especially the second factor is an undeniable problem for patients with Huntington's disease.

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 $b_h$.[18] 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 Proposition 2.* Applying Bayes' rule, we immediately obtain the requirements

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

and

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

Solving for $\alpha$ and $\beta$ yields the solution given in the proposition. It remains to check that $\alpha, \beta \in [0,1]$. For $\beta$ 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 3.* Since the optimal test is invariant to multiplying $V$ by a constant, we can reinterpret decreasing $\theta$ 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:

[18] 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.

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 4.* 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 $E[V(D(d_h))] = g_{d_h}(p)$ 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 weakly greater than $V$ over $[p,1]$. Clearly, $g(p) \geq E[V(D(d_h))]$ 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 $D(d_h^*)$ 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 Lemma 2.* We have to show that the second derivative $V''(b) = \theta U''(b) - (1-\theta)c^{*\prime\prime}(b)$ switches signs at most once and if it does then from negative to positive. Under our assumption on $C(x,\cdot)$, the function $c^*(b)$ is given by $c^*(b) = b(1-b)$ and thus $c^{*\prime\prime}(b) = -2$ for all $b$. Since $U''$ is monotone, $\theta U''(b)$ and $(1-\theta)c^{*\prime\prime}(b)$ intersect at most once and it is easily checked that the resulting signs of $V''$ match the claims in the lemma, and that the point of intersection $b_c$ moves to the right as $\theta$ increases. □

*Proof of Proposition 5.* 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 $b > b_t$. □

*Proof of Proposition 6.* 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 7.* Note first that the signals relevant to our restricted optimization are exactly those signals which induce beliefs $B \in \mathcal{B}_G$ where $\mathcal{B}_G \subset \mathcal{B}$ is the set of $G$-valued random variables with mean $p$. The requirement $\min G \leq p \leq \max G$ ensures that $\mathcal{B}_G$ is non-empty. Since $\overline{V}_G$ is concave and dominates $V$ on $G$, we have by Jensen's inequality the upper bound

$$\sup_{B \in \mathcal{B}_G} E\left[V(B)\right] \leq E\left[V_G(B)\right] \leq V_G(E[B]) = V_G(p)$$

on what the doctor can achieve. Choose $b_l$ and $b_h$ as prescribed in the proposition. Clearly, $b_l$ and $b_h$ must lie in $G$, since – due to its minimality – $V_G$ can only change its slope at points in $G$ and since the interval $(b_l, b_h)$ was chosen maximally. Thus, the random variable $B_G^*$ which takes only values $b_l$ and $b_h$ and has mean $p$ lies in $\mathcal{B}_G$. Since $V_G$ is linear over $[b_l, b_h]$, the optimum is attained: $E[V(B_G^*)] = V_G(p)$. □

*Proof of Corollary 2.* The proof is basically contained in the one of Proposition 4: In order to construct a test which induces a belief which is preferable to $D(b_h)$ we would need the existence of a straight line $g_{d_h}$ going through $(d_l, V(d_l))$ and some point $(d_h, V(d_h))$, $d_h \geq b_h$, with the property that $g_{d_h}$ is steeper than the line $g_{b_h}$ which connects $(d_l, V(d_l))$ and $(b_h, V(b_h))$. Yet we have shown that $g_{b_h}$ is strictly greater than $V$ on $(b_h, 1]$ and thus no steeper line through $(d_l, V(d_l))$ can intersect $V$ on $(b_h, 1]$. This proves the claim. □

*Proof of Proposition 8.* Recall first that the same values $b_l$ and $b_h$ characterize the optimal test for all priors $p' \in [b_l, b_h]$ and thus as long as the patient holds a belief within this interval the optimal test induces these values. By the definition of $b_l$ and $b_h$ in Proposition 1, we know that in $b_l$ and $b_h$ the functions $V$ and $\overline{V}$ coincide so that a patient who holds a posterior belief in $\{b_l, b_h\}$ will not accept any further offered tests. In particular, if $b_l = p = b_h$ the patient will never accept any test and the result follows trivially. Assume thus $b_l < p < b_h$. It thus remains to show that one of these beliefs is always reached through our procedure. The key observation is that if the patient has accepted a test with values in $\{d_l, d_h\}$ in our sequence of tests then we can conclude that $[d_l, d_h] \subseteq [b_l, b_h]$. Suppose that the starting interval for the current round of offers $[s_l, s_h]$ is contained in $[b_l, b_h]$ and that the current prior $p$ lies strictly between $b_l$ and $b_h$. Observe that in our sequence of tests, we always increase $d_h$ while keeping $d_l$ fixed until we reach $d_h = 1$ or the patient accepts a test. By Corollary 2 we know that if the patient does not accept the test with $\{d_l, b_h\}$ he will not accept any test with $\{d_l, d_h\}$, $d_h > b_h$ either. Yet since we decrease $d_l$ gradually and since the test $\{d_l, b_h\}$ is always offered before any too informative test with $\{d_l, d_h\}$, $d_h > b_h$, it is clear that if the patient accepts a test we know that it cannot include a value of $d$ outside $[b_l, b_h]$ – otherwise the patient would have accepted an earlier test. At the very beginning we have $\{s_l, s_h\} = \{p - \varepsilon, p + \varepsilon\}$ and thus $[s_l, s_h] \subseteq [b_l, b_h]$. Thus it follows inductively that the results $\{d_l, d_h\}$ of an accepted test always lie in $[b_l, b_h]$ and accordingly, the starting points for searching the next test always lie in $[b_l, b_h]$ as well. The only exception is the case where the patient has a posterior belief in $\{b_l, b_h\}$ after the test so that either $s_l = b_l - \varepsilon$ or $s_h = b_h + \varepsilon$. Yet this case is not problematic since a patient who holds an optimal belief refuses all further tests. Finally, it is clear that a posterior belief in $\{b_l, b_h\}$ is ultimately reached since a test inducing these two values is eventually proposed to the patient unless he has reached either of these values earlier. This follows immediately from the way the sequence of tests is constructed. □

*Proof of Lemma 3.* Let $\rho \in [0, 1]$ and $a, b \in \Delta_n$. Then concavity of $U$ follows from concavity of $u$ via

$$
\begin{aligned}
U(\rho a + (1-\rho)b) &= u\left(\rho\left(\sum_{i=1}^{n} a_i x_i\right) + (1-\rho)\left(b_i x_i\right)\right) \\
&\leq \rho\, u\left(\sum_{i=1}^{n} a_i x_i\right) + (1-\rho)u\left(\sum_{i=1}^{n} b_i x_i\right) \\
&\leq \rho\, U(a) + (1-\rho)U(b).
\end{aligned}
$$

□

*Proof of Proposition 9.* Recall that the optimal test takes the same two values $b_l(s)$ and $b_h(s)$ for all priors $p$ in an open interval around $p_0$ on which $V_s$ is strictly smaller than $\overline{V}_s$. Define $b_l(s)$ and $b_h(s)$ this way. Since we have $p_0 \in [b_l(s-1), b_h(s-1)] \subseteq [b_l(s), b_h(s)]$ by Proposition 3, the result on $b_l(s)$ and $b_h(s)$ follows immediately. Accordingly the test at time $s$ induces either $p(s+1) = b_l(s)$ or $p(s+1) = b_h(s)$ regardless of the value of $p(s)$. □

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

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

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

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

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

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

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

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

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

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

[13] Kőszegi, Botond. Utility from Anticipation and Personal Equilibrium, Economic Theory, 44, 415-444, 2010.

[14] Leland, Hayne E. Savings and Uncertainty: The Precautionary Demand for Saving, Quarterly Journal of Economics, 82, 465-73, 1968.

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

[16] Oster, Emily, Ray Dorsey and Ira Shoulson. Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease, American Economic Review, forthcoming, 2012.

[17] Rosar, Frank, and Elisabeth Schulte. Optimal test design under imperfect private information and voluntary participation, Working Paper, 2012.

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

[19] 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.