



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Health Economics 22 (2003) 1073–1084

JOURNAL OF
HEALTH
ECONOMICS

www.elsevier.com/locate/econbase

Health anxiety and patient behavior

Botond Kőszegi*

*Department of Economics, University of California, Berkeley, 549 Evans Hall #3880,
Berkeley, CA 94720, USA*

Received 1 August 2002; received in revised form 1 June 2003; accepted 20 June 2003

Abstract

Economic models of patient decision-making emphasize the costs of getting medical attention and the improved physical health that results from it. This note builds a model of patient decision-making when fears or anxiety about the future—captured as beliefs about next period's state of health—also enter the patient's utility function. Anxiety can lead the patient to avoid doctor's visits or other easily available information about her health. However, this avoidance cannot take any form: she will never avoid the doctor with small problems, and under regularity conditions she will never go to a *bad* doctor to limit the information received.

© 2003 Elsevier B.V. All rights reserved.

JEL classification: C7; D82; I19

Keywords: Health anxiety; Behavior; Economic model

1. Introduction

Economic models of patient decision-making emphasize the costs of getting medical attention and the improved physical health that results from it. This is what underlies many of the cost–benefit calculations that compare the value people attach to better health with the costs of medical intervention (Cutler, 2000, for example). Yet medical professionals and patients themselves are often also concerned about the stress patients must undergo when they face a serious health situation. Such stress, based largely on fears about future health, is not only unpleasant, but is thought to influence measurable health outcomes and patient behavior.

This note builds a theoretical model of patients' anxiety about their health and its consequences for their behavior. The analysis is motivated by medical evidence (reviewed in

* Corresponding author. Fax: +1-510-642-6615.

E-mail address: botond@econ.berkeley.edu (B. Kőszegi).

Section 2) that people often decline to learn free information about their health, and are reluctant to see a doctor with suspicious symptoms. I argue that this behavior is plausibly driven by fears of bad news. Then, following [Caplin and Leahy \(1999\)](#), I model a patient's fears as arising from her expectations about her future physical condition. Thus, in my model, the patient's utility function is defined not over physical outcomes as in standard analysis, but instead over *beliefs* about physical outcomes. The patient can learn free information about her condition, which she can use to choose a more appropriate treatment to undergo or a more fitting lifestyle to adopt. The note explores the types of information avoidance that can occur under these circumstances.

Since the patient derives utility directly from her beliefs, she must consider how the information she gathers will affect those beliefs. If she dislikes bad news more than she likes good news (a preference called information aversion), she trades off her fear of bad news with the decision-making value of knowing her state of health. As a result, she might refuse useful information that is very cheap. For example, she might be reluctant to conduct regular self-examinations or go for physicals, even though these are crucial in prevention. At first sight, one might expect that if the medical issue at hand is sufficiently serious, she will choose to fight her fears and learn any cheaply available information. Surprisingly, the opposite is the case in some sense. The note proves that the decision-maker will (almost) never avoid the doctor if the visit is useful and she expects to learn little from it, but may do so if very bad news are possible.

In addition, even though patients might not want to go to the doctor, under regularity conditions (and holding financial costs constant) their anxiety will never induce them to go to a *bad* doctor, whose information about their state of health is worse. In other words, the patient goes to the best doctor or no doctor at all. The intuition for this result is that if her aversion to information is weak enough so that she finds a doctor's advice valuable, she will find further information—which the better doctor can provide—valuable as well.

My model is based on [Caplin and Leahy's \(2001\)](#) psychological expected utility model. As I have mentioned, the same authors also apply a version of their model to study doctor–patient interaction ([Caplin and Leahy, 1999](#)). However, their paper has a completely different focus: while they deal with the strategic problems that arise in communication between a doctor and a patient, the current note deals exclusively with patient decision-making.

The paper is organized as follows. [Section 2](#) presents some medical evidence to motivate the model. [Section 3](#) introduces the utility function incorporating anxiety and the decision-making problem. [Section 4](#) derives the consequences of anxiety for individual decision-making. The final section concludes.

2. Motivating facts

This section presents a set of motivating facts from the medical literature on patient behavior. I argue that these facts are unlikely to be explained in a convincing way by standard models, and that a very plausible underlying motivation behind them is patients' emotional reactions to news about their health. However, I claim neither that there is incontrovertible evidence that anxiety affects medical decision-making, nor that any of the findings below support the more specific assumptions or predictions of my model.

In a series of studies, Lerman and colleagues have examined patients' preferences for information regarding their susceptibility to various forms of cancer. Specifically, Lerman et al. (1996b) find that 40% of *high-risk* patients who are offered a test for genetic susceptibility to breast and ovarian cancer declined the test. In a similar study on a type of colon cancer, 57% of high-risk individuals declined to know the test results (Lerman et al., 1999). Since the results were immediately and freely available to these patients, and their anonymity was guaranteed, there seems to be no rational or strategic reason to refuse this information. Lerman et al. (1996a) cite anticipated emotional reactions to bad news as one of the main barriers to testing.

Relatedly, there is ample (survey-type) evidence that patients often delay seeing a professional with symptoms they either consider suspicious or know to be serious. In a meta-analysis of 12 studies, Facione (1993) concludes that 34% of women with breast cancer symptoms delay help seeking for three or more months. Although some women seem to delay because they do not consider their symptoms to be serious, others know they should see a physician but are afraid, or hope for the symptoms to go away by themselves (Nosarti et al., 2000). Richard et al. (2000) find that the median melanoma patient delays seeing a doctor for two months after realizing that she has a suspicious lesion. Given that early detection is crucial in breast-, skin-, and other types of cancer, it is difficult to calibrate a standard model to explain this type of behavior: in order for the patient to risk her life by delaying, the added cost of visiting the doctor early rather than later must be enormous. In addition, since these studies rely on surveying patients who eventually decided to see a doctor, and many of them may not want to admit that they made a mistake, the findings might understate the extent of the problem.

Even more puzzling is evidence indicating that patients who seem to have more to gain from visiting the doctor are sometimes less likely to go. For example, Caplan (1995) reports that women with breast cancer symptoms which are getting worse delay *longer* in seeing a professional than those whose symptoms are steady or disappearing. And in Meechan et al. (2002), women who had a family member with breast cancer tended to delay longer in seeing a physician than those who did not. Once again, fears and anxiety offer a plausible explanation for all the above phenomena. The fact that self-reported cancer worry (Bowen et al., 1999) and psychological stress (Kash et al., 1992) is negatively and significantly related to past and subsequent participation in risk counseling lend some support to this hypothesis.

3. Setup

I first introduce the patient's utility function, which is different from those in standard models of patient behavior in a way that is crucial for all results in the note. There are two periods, 1 and 2, and all utility is ultimately derived from health outcomes in period 2. In period 2, the patient's state of health depends on her diagnosis in period 1, denoted s (for "symptoms"), and some action t taken in period 1. Depending on the context of the decision-maker's problem, t can mean at least two distinct things: it could be a treatment received from the doctor or a health-relevant lifestyle choice the patient makes. For example, for a woman who notices a lump in her breast, the diagnosis s can capture whether the lump

is malignant, and t could be the set of treatments available in such a situation: mastectomy, less invasive surgery, doing nothing, and so on.

Both variables s and t are one-dimensional: s can take any value in the interval (A, B) , with probability distribution function $f(s)$ and cumulative distribution function $F(s)$. A and B can be infinite. The treatment or lifestyle chosen, t , can take any value in the compact set $T \subset \mathbb{R}$. The patient's level of health in period 2 is then $s - l(s, t)$, with $l(s, t) \geq 0$ for all s, t , l continuous, and $\min_{t \in T} l(s, t) = 0$. The interpretation of the function $s - l(s, t)$ is the following: depending on the patient's luck (s), there is some maximum level of health she can achieve if the appropriate action is taken. If her treatment or lifestyle is not appropriate for her symptoms, she can be worse off. This is captured by the loss function $l(s, t)$.

One important property of this setup is that the highest attainable health level varies with the diagnosis s . That is, if s happens to be low, even the optimal treatment cannot raise the patient's health back to a level she could achieve if the diagnosis was favorable. This is a natural assumption for many medical situations, including the motivating examples in the previous section. If the patient is diagnosed with cancer, her future well-being is likely to be lower than if she was found to be healthy, even if she gets the best possible care. However, not all medical conditions satisfy this property, and I will highlight how the results would be different in that case.

Since the analysis will center around behavior in period 1, we need to specify how the patient incorporates her future health into her first-period utility function. I assume that the patient's first-period utility derives from her anticipation of health in the second period. Naturally, this anticipation can only depend on what she knows at the time. Specifically, her first-period von Neumann-Morgenstern utility function takes the form

$$u(E[s - l(s, t) | \text{patient's information}]) \quad (1)$$

The patient's utility depends on her expected health in period 2 conditional on her information, which she controls to an extent described below. I assume that there is some $\alpha > 0$ such that for any $x, y \in \mathbb{R}$ with $x > y$, we have $u(x) - u(y) \geq \alpha(x - y)$; this means that u is increasing, and the extent to which improvements in health increase utility is bounded away from zero. The shape of u determines the patient's preferences for information. If u is concave, she is called (analogously to risk aversion) "information-averse;" if u is convex, she is "information-loving;" and if u is linear, she is "information-neutral."¹ Finally, the patient has correct priors about s , interprets information correctly, and is an expected utility maximizer (she maximizes the expectation of u in any decision-making problem).

It might look strange to assume that the decision-maker cares only about her emotions in period 1. A more plausible starting point is that she cares about both anticipatory utility in period 1 and actual health in period 2. But such a model would generate identical behavior to one of the above type, because the patient's concern for future health can be included in u .

This model is an example of Caplin and Leahy's psychological expected utility model (Caplin and Leahy, 2001), and is in the spirit of Kreps and Porteus (1978). See Caplin and

¹ By the law of iterated expectations, the patient's average beliefs are unchanged by information. Therefore, by Jensen's inequality, her expected utility decreases if u is concave, and increases if u is convex.

Leahy for a careful defense of this variant of the expected utility approach, in which the substitution axiom is maintained for preferences over beliefs.

The patient's decision-making problem is the following. In period 1, she decides whether to costlessly learn about s or to remain uninformed. I first assume that in case she decides to learn s , she does so perfectly, although I will later consider the patient's preferences over *how* informed she wants to become. After learning s , the patient chooses an action t . Once again, this assumption attempts to capture multiple situations. In one example, the patient might or might not decide to perform a self-examination that identifies potential health problems. In either case, she has to decide what kind of health-related lifestyle (diet, exercise, etc.) to adopt. Another example is the patient's decision of whether to visit a doctor. In case she visits, the doctor gives her the diagnosis s , which she otherwise does not find out.² A treatment can then be chosen based on this knowledge.³

4. A few implications of anxiety

Whether or not the patient decides to learn s , she chooses the appropriate treatment *given her beliefs*—this is what maximizes expected health in the second period, and therefore also her anticipatory utility.⁴ The interesting implications of the model derive from the patient's desire to remain ignorant in some cases.

Let t^* denote the optimal treatment choice when the patient does not know s : $t^* = \operatorname{argmin}_t E[l(s, t)]$. t^* exists because l is continuous and T (the set from which t is chosen) is compact. If she does not learn s , her expected utility is thus $u(E[s - l(s, t^*)])$. If she learns s , she chooses $t = s$ in each state of the world, giving her an expected utility of $E[u(s)]$. The patient prefers not to learn her state of health if the former expression is greater than the latter. This condition can be written as

$$\underbrace{u(E[s]) - E[u(s)]}_{\text{info preferences}} > \underbrace{u(E[s]) - u(E[s - l(s, t^*)])}_{\text{improved treatment}}. \quad (2)$$

Eq. (2) summarizes the tradeoffs facing the patient. Learning s affects her knowledge about the future, and, depending on her informational preferences, this will in general change her utility from anticipation. This consideration is shown on the left-hand side of the inequality. On the other hand, learning s allows the patient to choose t optimally. This increases her health in the second period, and, in anticipation of the increased health, also makes her feel better in the first period. This effect is shown on the right-hand side of Eq. (2).

² Clearly, a doctor generally has superior information relative to a patient. Thus, it is in the doctor's power to hide at least some information from the patient, making the revelation of s a strategic decision. This consideration is ignored in the current note, but is the explicit topic of Caplin and Leahy (1999) and Kőszegi (2001).

³ In my model, the set of treatments available to the patient is independent of whether she learns s . Clearly, this is not always the case; for example, one cannot be treated for cancer without first being diagnosed with it. Including this consideration would not affect the results.

⁴ In the example of visiting the doctor, the patient would instruct the doctor to give her the treatment $t = s$, instead of making the lifestyle choice herself. A doctor who maximizes the patient's utility from health would make the same decision.

The right-hand side of Eq. (2) is always non-negative. Thus, if the left-hand side is negative or zero, the patient prefers to learn s . The left-hand side is negative if the patient is information-loving (u is convex), and it is zero if she is information-neutral (u is linear). Naturally, if the patient does not dislike information, the fact that learning the diagnosis helps her in choosing treatments is sufficient for her to want to know s .

However, if our patient is information-averse (u is concave), she is weighing off feeling better today through ignorance and being healthier tomorrow. If the former consideration outweighs the latter, the patient will choose not to learn s . As the observation below demonstrates, either effect could dominate. In the statement of the observation, $\text{supp}(F)$ denotes the support of the distribution F .

Observation 1. *If $E[s - l(s, t^*)] \leq \min \text{supp}(F)$, then the patient chooses to learn s for any u . If $E[s - l(s, t^*)] > \min \text{supp}(F)$, then there is a u for which the patient chooses not to learn s .*

Proof. Appendix A. □

If the treatment chosen without knowing the precise diagnosis is so bad that the average level of health is lower than the worst possible diagnosis treated properly ($E[s - l(s, t^*)] \leq \min \text{supp}(F)$), then the patient will choose to learn s , no matter how afraid she is of bad news. The reason is simple: even if she received the worst possible news, her anticipatory utility would still be higher than if she did not know s , but knew she is choosing the treatment suboptimally. However, if choosing t^* in ignorance of s does not yield such an extremely low expected utility, the patient may prefer not to know s . This happens if she is very afraid of the very worst news.

As some of the empirical evidence indicates, patients are often information-averse. In addition, there is also a good theoretical reason to expect decision-makers to be information-averse a lot of the time. As shown in Kőszegi (2000, Section 3), if the patient has an abundance of voluntary opportunities to learn about her health, she keeps learning as long as she is on an information-loving part of her utility function.⁵ Thus, when she stops learning, she is more likely to be on an information-averse part.

Two assumptions are crucial in generating the model's prediction that the patient might choose not to learn s . First, as mentioned above, I assumed that the diagnosis affects the maximum utility the patient can achieve with the optimal treatment. For some medical conditions, this is clearly not the case, because the available treatments bring the patient back to perfect health. For example, if the patient has diarrhea, this is easily curable whether it is food poisoning or a stomach virus.

To model this alternative possibility, I assume that the patient's level of health in the second period is $-l(s, t)$ instead of $s - l(s, t)$, so that her maximum level of utility is the same (and normalized to zero) for any s . As before, I assume that anticipatory utility

⁵ Kőszegi (2000) deals with ego utility, the notion that people like to see themselves as capable individuals. This seems different from anxiety. However, ego utility turns out to be quite similar to anxiety theoretically, because both can be modeled as utility from beliefs. This allows me to apply the previous paper to a framework with anxiety.

depends on expected health in the second period conditional on the patient’s information. Then, the left-hand side of Eq. (2) becomes $u(E[0]) - E[u(0)] = 0$, so the patient always chooses to learn s . Intuitively, since any possible diagnosis is perfectly treatable, there is no sense in which some diagnoses are bad news (or good news), so the patient’s informational preferences do not affect her choice.

Second, the model above assumes that the patient needs to know s to choose the correct t . Once again, this is a reasonable assumption for the examples—such as cancer screening—used to motivate the model. For other kinds of preventive care, however, the patient does not need to receive any new information to take the right action. Washing hands, putting on sunscreen, and wearing hard hats at construction sites are examples. In my model, anxiety cannot lead the patient to underuse such information-insensitive precautions.

Besides showing conditions under which the patient might choose not to undergo useful screening and diagnostic procedures, this model can say more about the forms this avoidance will take. The rest of this section derives two results on the limits of information-avoidance behaviors.

The first result has to do with the patient’s attitude towards learning “small news,” news that are not going to move her beliefs by much. To model this, I take a distribution of diagnoses with mean \bar{s} , “shrink” its support toward \bar{s} —so that eventually most of the weight is around \bar{s} —and see how the patient behaves in the limit. When the distribution is concentrated around \bar{s} , learning s will not affect the patient’s anxiety much, making the potential utility loss due to anticipatory utility small. On the other hand, the treatment chosen in full knowledge of s will not be much better than the one chosen based on the prior distribution. Although both sides of Eq. (2) thus approach zero as the distribution of s shrinks toward \bar{s} , the patient will tend to want to learn s :

Theorem 1. *Suppose that there is a constant $k > 0$ such that $l(s, t) \geq k |s - t|$ for all s, t . Take any continuous probability distribution g over the interval $[a, b]$ with mean zero. For any $\epsilon > 0$, define*

$$f_{\bar{s}, \epsilon}(s) = \begin{cases} \frac{1}{\epsilon} g\left(\frac{s - \bar{s}}{\epsilon}\right), & \text{if } s \in [\bar{s} + \epsilon a, \bar{s} + \epsilon b] \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

For almost all $\bar{s} \in (A, B)$, if ϵ is sufficiently small, a patient facing a distribution of diagnoses $f_{\bar{s}, \epsilon}$ will choose to learn s .

Proof. Appendix A. □

Therefore, anxiety can keep patients away from seeking help only in serious cases, exactly the opposite of what standard economic logic would predict! In the usual cost-benefit calculus, a more serious symptom implies that there is more to gain from appropriate care, so patients will be more likely to visit the doctor. However, the anxiety costs of getting care can also increase with the seriousness of the symptoms, and it turns out that they can only outweigh the benefits in more serious cases. For example, a woman might be too afraid to go to the doctor when she notices a lump in her breast, but at the same time have no qualms about going with an equally annoying inflammation, in which case cancer is very unlikely.

Theorem 1 is a consequence of the expected utility formulation of anxiety, in which the patient maximizes the expectation of a utility function that depends on beliefs. Since u is increasing, it is differentiable almost everywhere, so for almost every initial belief, it is linear locally. But a linear u implies information neutrality, so the patient is “locally information-neutral.” Intuitively, since the patient’s beliefs are not going to move to a much steeper part of u , she is just as likely to feel better as a result of the news as worse. If the gains from going to the doctor and getting diagnosed are first-order—that is, knowing s accurately still improves the choice of t in the limit—the patient will therefore go to the doctor.⁶

Theorem 1 applies to situations where there is little uncertainty about the patient’s health, and the Theorem says that in this case the patient is willing to learn her state of health. Importantly, this does *not* imply that the patient will want to find out at least a little bit about a large amount of uncertainty. (If that was the case, she would always want all information, since she would never be satisfied with any amount of partial learning.) The crucial distinction is that adjustments in the choice of t as a result of partial improvements in information are not as valuable as those that result from full learning, because in the former case the adjustment could go the wrong way.

The model’s implication that the patient avoids a doctor’s visit that could potentially improve her health raises an intriguing possibility: would people prefer to visit a *bad* doctor rather than a good one to limit information about their health? A major aspect in which a bad doctor differs from a good one is that she observes the diagnosis s with more noise, decreasing what the patient can learn from her. Since holding costs are constant, it seems unlikely that patients would willingly choose bad doctors, a model that generates this kind of behavior is questionable. Fortunately, under reasonable regularity assumptions, a patient will either visit no doctor at all, or will go to the best doctor she can find, as **Theorem 2** demonstrates.

Theorem 2. *Suppose that $u(H) = e^{-rH}$, $l(s, t) = |s - t|$, s is normally distributed, and that only a noisy version of s , $s_d = s + \epsilon$ is available, where $\epsilon \sim N(0, \sigma_d^2)$. If the patient is willing to observe a noisy diagnosis with variance σ_d^2 she prefers any diagnosis with lower variance.*

Proof. Appendix A. □

The intuition behind **Theorem 2** is related to the previous result that the patient is almost always willing to find out small pieces of information. Since utility is a concave function of expectations, small amounts of information are relatively more valuable to the patient than big ones (they are less likely to move the patient to a very steep part of u). This, then, implies that if she is willing to go to some doctor (a large amount of information), she would rather go to a slightly better doctor (which is a small amount of extra information). And since this

⁶ This implication of the independence axiom of expected utility theory is a very intuitive one. As demonstrated by the Ellsberg (1961) and Allais (1953) paradoxes, for example, many other implications of the axiom are less reasonable. What makes the prediction in **Theorem 1** more intuitive is probably its greater realism about the *payoffs* that enter expected utility. In the standard setup, decisionmakers are assumed to care only about physical payoffs like monetary rewards, essentially a misspecification of the model. If the domain of utility is made more realistic, the predictions of the theory might also be more in line with observed behavior.

is true for any doctor, she prefers the best doctor.⁷ For this logic to work, two conditions need to be satisfied. First, the value of information for choosing t cannot vary over the range of possible diagnoses (l is linear in $|s - t|$). For example, if l was a function of $|s - t|$, and was flat for small $|s - t|$, but then increased sharply, it may be important to narrow down s sufficiently, with further information being less useful. Second, the patient's information aversion cannot vary over the range of her beliefs (u is exponential). Otherwise, the patient could be information neutral locally, but be extremely information averse further away, and as a consequence prefer to receive some information, but not a lot of information.

5. Conclusion

This paper provides the first economic model of patient decision-making when anxiety affects choices regarding information acquisition and treatment. I model anxiety by assuming that the agent derives utility from her beliefs about her future health. Even without making many assumptions about the functional form anticipatory feelings might take, the model can deliver strong and intuitive results about patient behavior.

Working closely with medical researchers, future work can perhaps put more structure on the utility function, allowing the development of more specific predictions. One logical next step is to study the effects of health education and the increase in the availability of medical information on the avoidance behaviors identified in this paper. If health education allows some patients to at least partially diagnose themselves, they may be more willing to go to the doctor, since such a visit is now less informative (but still useful for treatment). But if health education makes doctors visits more informative—because the patient can make more of a doctor's diagnosis—it can actually exacerbate the avoidance problem.

More importantly, the role of doctors, and the medical establishment in general, is taken to be rather limited in this model. The doctor simply functions as a black box, as an information provider making no decision of her own. In reality, doctors make many strategic decisions, including what to tell their patients and what to leave out. This complicates the model considerably, resulting in a signaling game that is taken up in [Caplin and Leahy \(1999\)](#) and [Kőszegi \(2001\)](#). In addition, the consideration of anticipatory emotions raises a host of questions regarding government policy that require further scrutiny. These include normative issues such as the extent to which a social planner should take feelings into account, as well prescriptive ones such as ways to encourage “correct” patient behavior.

Appendix A. Proofs

Proof of Observation 1. If $E[s - l(s, t^*)] \leq \min \text{supp}(F)$, then $u(E[s - l(s, t^*)]) \leq u(\min \text{supp}(F)) \leq E[u(s)]$ so the patient prefers to learn s .

⁷ Though much less obvious, [Theorem 2](#) is reminiscent of Samuelson's result that if a decision-maker accepts one-hundred identical positive expected value gambles, she should accept a single one as well ([Samuelson, 1963](#)). Going to a doctor is a “big” gamble on beliefs, whereas going to a slightly better one is a “small” extra one. As in Samuelson's case, if the patient takes the big gamble, she should also take the small one.

Now suppose that $x \equiv E[s - l(s, t^*)] > \min \text{supp}(F)$. Since $\text{Prob}(s < x) > 0$, $u(x) - \text{Prob}(s < x)E[u(s)|s < x]$ can be arbitrarily large, with u remaining bounded from above. (More precisely, there is a real M such that for any real K , there is a u that is bounded from above by M and $u(x) - \text{Prob}(s < x)E[u(s)|s < x] > K$. Thus, $u(x) - E[u(s)]$ can be positive, in which case the patient does not want to learn s . \square

Proof of Theorem 1. We prove that the statement is true for any \bar{s} where u is differentiable. To do this, we once again use Eq. (2), and prove that for a sufficiently small ϵ , the right hand side is greater than the left-hand side. Notice that

$$E_{f_{\bar{s},\epsilon}}[l(s, t^*)] \geq E_{f_{\bar{s},\epsilon}}k|s - t^*| \geq \min_{t \in T} E_{f_{\bar{s},\epsilon}}k|s - t|.$$

Through a change of variables, the above equals

$$\epsilon k \min_{t \in T} E_g \left| s - \frac{t - \bar{s}}{\epsilon} \right| \geq \epsilon k \min_{t \in \mathbb{R}} E_g \left| s - \frac{t - \bar{s}}{\epsilon} \right| \equiv \epsilon K,$$

where K is a constant. Thus, the right-hand side of Eq. (2) can be bounded from below:

$$u(E[s]) - u(E[s - l(s, t^*)]) \geq \alpha E_{f_{\bar{s},\epsilon}}[l(s, t^*)] \geq \epsilon \alpha K.$$

On the other hand, $1/\epsilon$ times the left-hand side of Eq. (2) is

$$\frac{u(\bar{s}) - E_{f_{\bar{s},\epsilon}}u(s)}{\epsilon} = -E_{f_{\bar{s},\epsilon}} \frac{u(s) - u(\bar{s})}{s - \bar{s}} \cdot \frac{s - \bar{s}}{\epsilon}. \tag{A.1}$$

Since u is differentiable at \bar{s} , the first fraction in the integrand approaches a constant (the derivative) as $\epsilon \rightarrow 0$. Since the second fraction is bounded, the limit of the expectation is equal to $E((s - \bar{s})/\epsilon)$, which is zero. Therefore, for a sufficiently small ϵ , the right-hand side of Eq. (A.1) is less than αK . As a consequence, $u(\bar{s}) - E_{f_{\bar{s},\epsilon}}u(s) < \epsilon \alpha K$, so the patient prefers to learn s . \square

Proof of Theorem 2. We prove that if the patient observes the signal with variance σ_d^2 then she will afterwards be willing to observe another signal of sufficiently small variance. This is sufficient to prove the result.

When the agent’s beliefs have mean μ and variance σ^2 her anticipatory utility is $u(\mu - |E[s - \mu]|)$ since she would choose $t = \mu$ in that case. Notice that $|E[s - \mu]|$ is constant in μ (i.e. it only depends on σ^2). Call this constant $k(\sigma)$. Thus, the agent’s anticipatory utility is $u(\mu - k(\sigma))$.

Now consider what happens when the agent gets a signal about s . As a result of observing a signal, the agent’s beliefs change; in particular, her posterior mean is a random variable that is normally distributed; suppose its variance is σ_m^2 . The variance of the agent’s beliefs conditional on the signal is then $\sigma^2 - \sigma_m^2$. The resulting loss function is $k(\sigma^2 - \sigma_m^2)$. Therefore, drawing the signal is a gamble on expected health with mean $k(\sigma^2) - k(\sigma^2 - \sigma_m^2)$ and variance σ_m^2 . Since u is the exponential utility function, the agent accepts this bet iff

$$\underbrace{k(\sigma^2) - k(\sigma^2 - \sigma_m^2)}_{\text{info gain}} - \underbrace{\frac{r}{2}\sigma_m^2}_{\text{anxiety loss}} \geq 0. \tag{A.2}$$

We continue by studying the two terms on the left-hand side of this inequality. The anxiety loss is just the expectation of $(r/2)x^2$, where $x \sim N(0, \sigma_m^2)$. The info gain is a bit more complicated. When the agent gets a signal that moves her mean beliefs from μ to μ' , the gain in expected health from choosing $t = \mu'$ instead of $t = \mu$ is clearly a function of $|\mu - \mu'|$. Let the gain be denoted by $I_{\sigma^2 - \sigma_m^2}(|\mu - \mu'|)$. It is easy to show that $I_{\sigma^2 - \sigma_m^2}(0) = 0$ and $I_{\sigma^2 - \sigma_m^2}(x)$ is a twice differentiable function with first derivative $2\Phi_{\sigma^2 - \sigma_m^2}(x) - 1$ and second derivative $2\phi_{\sigma^2 - \sigma_m^2}(x)$, where $\Phi_{\sigma^2 - \sigma_m^2}$ and $\phi_{\sigma^2 - \sigma_m^2}$ are the cumulative and probability distribution functions of a normal with mean zero and variance $\sigma^2 - \sigma_m^2$, respectively. Thus, $I_{\sigma^2 - \sigma_m^2}(x)$ is convex with a second derivative that decreases in x . Finally, for any x , $I_{\sigma^2 - \sigma_m^2}(x)$ is non-increasing in the posterior variance $\sigma^2 - \sigma_m^2$.

Now notice that the info gain is the expectation of $I_{\sigma^2 - \sigma_m^2}(x)$ when $x \sim N(0, \sigma_m^2)$. Suppose then that inequality (A.2) is satisfied, that is, that the patient would want to learn this piece of information about s . By the properties of $I_{\sigma^2 - \sigma_m^2}$ for its expectation to outweigh that of $\frac{r}{2}x^2$, it must be greater around zero (otherwise it would be smaller everywhere.) This implies that for any $\sigma_p^2 < \sigma^2 - \sigma_m^2$, $I_{\sigma_p^2}$ is also greater around zero than $(r/2)x^2$. Thus, for a distribution sufficiently concentrated around zero, the expectation of $I_{\sigma_p^2}$ is greater than the expectation of $(r/2)x^2$. Thus, the agent is willing to take a sufficiently small extra risk on her beliefs. \square

References

- Allais, M., 1953. Rational behavior in the face of risk: a critique of certain postulates and axioms of the American School. *Econometrica* 21, 503–505.
- Bowen, D., Mctiernan, A., Burke, W., Powers, D., Pruski, J., Durfy, S., Gralow, J., Malone, K., 1999. Participation in breast cancer risk counseling among women with a family history. *Cancer, Epidemiology, Biomarkers, and Prevention* 8, 581–585.
- Caplan, L.S., 1995. Patient delay in seeking help for potential breast cancer. *Public Health Reviews* 23 (3), 263–274.
- Caplin, A., Leahy, J. 1999. The supply of information by a concerned expert. C.V. Starr Center for Applied Economics Working Paper.
- Caplin, A., Leahy, J., 2001. Psychological expected utility and anticipatory feelings. *Quarterly Journal of Economics* 116 (1), 55–80.
- Cutler, D.M., 2000. Does the United States spend too much on medical care? In: Garber, A. (Ed.), *Health Care in the United States and Japan*, forthcoming.
- Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75 (4), 643–669.
- Facione, N.C., 1993. Delay versus help seeking for breast cancer symptoms: a critical review of the literature on patient and provider delay. *Social Science and Medicine* 36 (12), 1521–1534.
- Kash, K.M., Holland, J.C., Halper, M.S., Miller, D.G., 1992. Psychological distress and surveillance behaviors of women with a family history of breast cancer. *Journal of the National Cancer Institute* 84 (1), 24–30.
- Kőszegi, B., 2000. Ego Utility, Overconfidence, and Task Choice, Mimeo.
- Kőszegi, B., 2001. Anxiety, Doctor–Patient Communication, and Health Education, Mimeo, UC Berkeley, URL: <http://emlab.berkeley.edu/users/botond/doctors.pdf>.
- Kreps, D.M., Porteus, E.L., 1978. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46 (1), 185–200.
- Lerman, C., Hughes, C., Trock, B.J., Myers, R.E., Main, D., Bonney, A., Abbaszadegan, M.R., Harty, A.E., Franklin, B.A., Lynch, J.F., Lynch, H.T., 1999. Genetic testing in families with hereditary nonpolyposis colon cancer. *Journal of the American Medical Association* 281 (17), 1618–1622.
- Lerman, C., Marshall, J., Audrain, J., Gomez-Camirero, A., 1996a. Genetic testing for colon cancer susceptibility: anticipated reactions of patients and challenges to providers. *International Journal of Cancer* 69 (1), 58–61.

- Lerman, C., Narod, S., Schulman, K., Hughes, C., Gomez-caminero, A., Bonney, G., Gold, K., Trock, B., Main, D., Lynch, J., Fulmore, C., Snyder, C., Lemon, S.J., Conway, T., Tonin, P., Lenoir, G., Lynch, H., 1996b. BRCA1 testing in families with hereditary breast-ovarian cancer. A prospective study of patient decision making and outcomes. *Journal of the American Medical Association* 275 (24), 1885–1892.
- Meehan, G., Collins, J., Petrie, K., 2002. Delay in Seeking Medical Care for Self-Detected Breast Symptoms in New Zealand Women. *The New Zealand Medical Journal* 115 (1166).
- Nosarti, C., Crayford, T., Roberts, J.V., Elias, E., Mckenzie, K., David, A.S., 2000. Delay in presentation of symptomatic referrals to a breast clinic: patient and system factors. *British Journal of Cancer* 82 (3), 742–748.
- Richard, M.A., Grob, J.J., Avril, M.F., Delaunay, M., Gouvernet, J., Wolkenstein, P., Souteyrand, P., Dreno, B., Bonerandi, J.J., Dalac, S., Machet, L., Guillaume, J.C., Chevrant-Breton, J., Vilmer, C., Aubin, F., Guillot, B., Beylot-Barry, M., Lok, C., Raison-Peyron, N., Chemaly, P., 2000. Delays in diagnosis and melanoma prognosis (i): the role of patients. *International Journal of Cancer* 89 (3), 271–279.
- Samuelson, P., 1963. Risk and uncertainty: a fallacy of large numbers. *Scientia* 98, 1–6.