Online Appendix for "Optimal Expectations and Limited Medical Testing"

Emily Oster, Ira Shoulson, E. Ray Dorsey

May 16, 2012

Appendix A: Continuous Actions

In this appendix we develop a version of the optimal expectations model with continuous action choices. We focus on the case of quadratic loss, and assume symmetry ($\Omega = \Phi$). That is, we assume that the loss to moving away from the correct action is the same regardless of the true state. Given that a central argument in favor of this model is that it works with this symmetry assumption, it seems reasonable to develop the continuous action case with this assumption. It is worth noting that an extreme form of asymmetry in which the loss to taking the wrong action if the true state is "sick" is much larger than the loss to taking the wrong action when the true state is "healthy" could cause this to break down. But this is the case in the binary model as well. The goal in this appendix is simply to demonstrate that the assumption that action a is binary is not necessary to derive our results.

Setup

We retain the setup from the main model. The state is binary, the timing is identical and, as there, individuals may choose beliefs π which differ from their true probability of carrying the HD mutation, which is p. We continue to denote utility from action a in state s as $u(a, s)$ and, as before, individuals choose time 1 beliefs $\pi \in [0, 1]$ to maximize

$$U(\pi|p) = \delta E(u(\hat{a}, s)|\pi) + E(u(\hat{a}, s)|p)$$

where $\hat{a}(\pi) = argmax_a E[u(a, s)|\pi]$. With the assumption of the binary state, we note that $E[u(\hat{a}, s)|p] = pu(\hat{a}, 1) + (1 - p)u(\hat{a}, 0)$.

We modify the model here to assume that action a is continuous: $a \in [0, 1]$, and that individuals experience quadratic losses associated with the difference between a and the true state. In particular, we write:

$$\begin{aligned} u(a, 1) &= -\Omega(1 - a)^2 \\ u(a, 0) &= 1 - \Omega(a)^2 \end{aligned}$$

As before, we assume that $\Omega < 1$.

Results

We develop our results as in the analysis in Section 4.2.

Appendix Proposition 1. Actions Given Beliefs $a = \pi$**.**

Proof. Actions are chosen in this model based only on the period 1 anticipation utility. Individuals will choose a to maximize

$$\pi[-\Omega(1-a)^2] + (1-\pi)[1-\Omega a^2]$$

Taking the first order condition and setting it equal to zero yields the result that $a = \pi$.

□

Action is matched to belief; the only difference from the binary case is that a will be directly matched to π if $\pi > 0$.

Appendix Proposition 1. Choice of Beliefs. *Individuals will always choose beliefs such that* $\pi \leq p$*. Beliefs* $\pi = 0$ *will be chosen for values of* $p^* \leq \frac{1+\delta\Omega}{2\Omega}$

Proof. Given that $a = \pi$ individuals will choose π to maximize

$$\delta[\pi(-\Omega(1-\pi)^2) + (1-\pi)(1-\Omega(\pi)^2))] + p(-\Omega(1-\pi)^2) + (1-p)(1-\Omega\pi^2)$$

This yields a closed form solution for $\pi : \pi = \frac{2p\Omega-\delta(\Omega+1)}{2\Omega(1-\delta)}$. Subtracting this from p we find that $p - \pi = \frac{-2\delta\Omega p+\delta(\Omega+1)}{2\Omega(1-\delta)}$. Since $\delta < 1$ the denominator is positive, and because $\Omega < 1$ the numerator is positive even for $p = 1$. Therefore $p > \pi$ for all p. In addition, we note that if the optimal $\pi < 0$ we will end up at a corner solution, and $\pi = 0$ will be chosen. This will occur if $\frac{\Omega(2p-\delta)-1}{2\Omega(1-\delta)} \leq 0$. Simplified, this implies that $\pi = 0$ will be chosen for values of $p^* \leq \frac{\delta(\Omega+1)}{2\Omega}$.

□

The choice of actions follows directly from the lemma and proposition. Individuals will choose action $a = 0$ up to values $p \leq p^*$ and then action $a = \pi$ for values of p above this. We note, of course, that it is easy to deliver values of $p^* \geq .5$, especially with high values of δ. For example, if we take the case developed in Figure 6, where $\delta = .8$ and $\Omega = .9$ we find that $p^* = .84$.

The result on testing and risk is virtually identical to the binary case. Individuals will be taking action $a = 0$ up to some p^*. The value of testing is increasing in risk over this range, and decreasing in risk above that value. We replicate the proposition here, with appropriate modification, but note that it is nearly identical to the binary case.

Appendix Proposition 2. Testing Behavior. *Define* p^* *as above. There are two cases, corresponding to a high and low anticipation value.*

Low Value of Anticipation: $\delta < \Omega$*. The following statements hold:*

1. *For values of* $p \leq p^*$*, the value of testing is positive if and only if* $p(\Omega - \delta) > C$
2. *For values of* $p > p^*$*, the value of testing is positive if and only if* $p(\Omega - \delta) - [\frac{(2p\Omega)^2(2\delta-1)-4\delta^2 p\Omega(1+\Omega)+\delta^2(1+\Omega)^2}{4\Omega(1-\delta)}] > C$

High Value of Anticipation: $\delta \geq \Omega$*. The following statements hold:*

1. *For values of* $p \leq p^*$*, the value of testing is negative and decreasing in* p*.*
2. *For values of* $p > p^*$*, the value of testing is positive if and only if* $p(\Omega - \delta) - [\frac{(2p\Omega)^2(2\delta-1)-4\delta^2 p\Omega(1+\Omega)+\delta^2(1+\Omega)^2}{4\Omega(1-\delta)}] > C$

Proof. If $p \leq p^*$, individuals take action $a = 0$. If $p > p^*$ they take action $a = \frac{2p\Omega - \delta(\Omega+1)}{2\Omega(1-\delta)}$. The value of testing for each range is given below.

$$V_{test} = p(\Omega - \delta) - C \text{ if } p \leq p^*$$
$$V_{test} = p(\Omega - \delta) - [\frac{(2p\Omega)^2(2\delta - 1) - 4\delta^2 p\Omega(1+\Omega) + \delta^2(1+\Omega)^2}{4\Omega(1-\delta)}] - C \text{ if } p > p^*$$

Low Value of Anticipation: $\delta < \Omega$. For $p \leq p^*$ and $a = 0$, the value of testing is increasing in p (since $\Omega - \delta > 0$), meaning it is maximized at p^*. For values of $p > p^*$ and action $a = 1$, the value of testing may or may not be increasing in p depending on parameter values. Putting this together, we note that testing is at least increasing in p up to $p^*$and it may be increasing in p above that, as well.

High Value of Anticipation: $\delta \geq \Omega$. For $p \leq p^*$, when $V_{test} = p(\Omega - \delta) - C$ it is easy to see that this value is negative as long as $\delta \geq \Omega$. It is also decreasing. Above a value of $p^*$the utility from testing may be increasing or decreasing in p.

□

Up to the critical value of p^*,the implications delivered here are the same as those in the binary case. In the binary model for higher values of p the value of testing was decreasing. Here, that may or may not be true; it is possible that the value of testing continues to increase above p^*. In this sense the testing and risk result may (under some parameter values) be even stronger in this continuous action, quadratic loss, case.

More broadly than the quadratic case, we note that the optimal expectations model as developed with general utility in Brunnermeier and Parker (2005) predicts that individuals will adopt overly optimistic beliefs and, as a result, their actions will be at least somewhat skewed. How skewed the actions are depends, of course, on functional form and parameter values. If anticipatory utility is sufficiently important, beliefs (and actions) may be very skewed. This follows directly from their paper. What is new here is the result on testing and risk. This appendix makes clear the broad implications are similar in the binary and continuous case. The fact that testing is increasing in risk comes from the fact that actions are sufficiently skewed that the maximized utility under the uninformed case grows more slowly than under the informed case.

A general condition for an arbitrary continuous loss function can be derived from the envelope theorem and is available from the author. Of course, as in the binary case, only some functional forms and parameter values will deliver the exact results we observe. This discussion serves only to provide a possibility result and to emphasize that the results in the paper are not special to the binary action setup.

Appendix B: Alternative Non-Neoclassical Models

In this section we describe two alternative, non-neoclassical models which have been suggested to explain these behaviors.

Wishful Thinking (Mayraz, 2011)[1]

The key assumption of this model is that individuals engage in "wishful thinking": when faced with a bad state of the world, they are overly optimistic. More concretely, individuals hold subjective beliefs which depend on the payoff function. We describe the setup and results below.

Setup

As before, there is a binary state $s \in \{0, 1\}$ and individuals are endowed with $p = E(s)$. The starting point in this model is bias in subjective beliefs: the assumption that individuals engage in wishful thinking causes their beliefs to be (potentially) biased toward the healthy state. As in the basic setup, we assume that individuals prefer to be healthy than to be sick. We denote the utility from health as $f(s)$. As in the body of the paper, we assume that $f(s = 1) = 0$ and $f(s = 0) = 1$, implying that the maximum possible utility in the healthy state is 1 and in the sick state is 0.

In general, Mayraz (2011) describes the formulation of subjective beliefs $\pi(s)$ from objective probabilities $p(s)$ as: $\pi(s) = p(s)e^{\psi f(s)}$ where ψ represents the level of optimism. Someone with a value of $\psi = 0$ is a realist; high values of ψ indicate high optimism. This model can also accommodate pessimism (with a negative ψ). In the HD case, the subjective probability can be computed

$$\pi = \frac{pe^{\psi(f(s=1))}}{pe^{\psi(f(s=1))} + (1-p)e^{\psi f(s=0)}} = \frac{p}{p + (1-p)e^{\psi}}$$

As in the baseline model, individuals choose a binary action $a \in \{0, 1\}$. Utility in this model is delivered in the same way as in the optimal expectations model:

$$U((a, s)|p) = \delta E(u(\hat{a}, s)|\pi) + E(u(\hat{a}, s)|p)$$

where $\hat{a} = argmax_a E[u(a, s)|\pi]$. However, in this model individuals consider *only* the anticipatory utility period when making choices about testing. In the optimal expectations model individuals make a choice of action based only on the beliefs in the anticipation period, but make choices about testing based on both periods. In this case, both action and testing choices are made based only on the anticipatory period, so we can write the object that individuals actually maximize as $U((a, s)|p) = E[u(\hat{a}, s)|\pi]$. For simplicity, we assume symmetric values for utility losses, so we have:

$$\begin{aligned} u(0, 1) &= -\Phi \\ u(1, 1) &= 0 \\ u(1, 0) &= 1 - \Phi \\ u(0, 0) &= 1 \end{aligned}$$

We assume that testing carries a real financial cost C.

Results

Results on beliefs and actions are shown in Appendix Proposition 3.

Appendix Proposition 3. Beliefs and Actions

1. *If $\psi > 0$ individuals hold beliefs $\pi < p$.*

[1] We are extremely grateful to Guy Mayraz for very helpful conversation about this model, and for developing the version of the model we present here.

2. *Action $a = 0$ will be taken for values of $p^* \leq \frac{e^\psi}{1+e^\psi}$*

Proof. 1. As noted above, in this model $\pi = \frac{p}{p+(1-p)e^\psi}$. As long as $\psi > 0$ this is less than p.

2. Lemma 1 in the paper describes the choice of action based on π. The model here corresponds to the symmetric case, so the cutoff value to switch from $a = 0$ to $a = 1$ is $\pi = .5$. This statement follows directly from solving for p when $\pi = .5$

□

The intuition for this result is similar to the optimal expectations case, although the result is delivered with one fewer step. By assuming biased beliefs, we deliver that result immediately, and the actions follow.

We next turn to testing behavior. The proposition below describes the relationship between testing and risk, the condition for individuals to test at all, and the condition for confirmatory testing. As in the optimal expectations case, we assume that testing for confirmation carries a value Ψ.

Appendix Proposition 4. Testing Behavior

1. *Value of testing is increasing in p for values of $p < p^*$ where $p^* = \frac{e^\psi}{1+e^\psi}$. Value of testing is decreasing in p for values of $p \geq p^*$.*

2. *Individuals will test if either:*

   (a) $p < p^*$ and $\pi(p)\Phi > C$ or

   (b) $p \geq p^*$ and $(1 - \pi(p))\Phi > C$

3. *Confirmatory testing will occur without predictive testing if and only if $\Psi > \Phi$.*

Proof. 1. Both points (1) and (2) follow from the value of testing. If $\pi(p) \leq .5$ and individuals are taking action $a = 0$ then value of testing is $V_{test} = \pi(p)\Phi$. If $\pi(p) > .5$ and individuals are therefore taking action $a = 1$ then $V_{test} = (1 - \pi(p))\Phi$. Note that p^* is the value such that $\pi(p^*) = .5$. Testing value is therefore increasing up to this $p^*$and decreasing for higher values. Moreover, in either case individuals will test if the value to testing exceeds the real cost.

2. To have confirmatory testing without predictive testing requires that individuals want to test for confirmation but not for predictive reasons. Testing for confirmation requires that $\Psi > C$. *Not* testing predicatively requires that the value to testing for confirmation only be higher than the value to testing predictively, where the latter is evaluated at the highest testing value point (p^*, or $\pi = .5$), at which point individuals are still taking action $a = 0$. These values are given:

$$V_{test,predictive} = .5u(1,1) + .5u(0,0) - C + \pi\Psi$$
$$V_{test,confirmation} = .5u(0,1) + .5u(0,0) + .5(\Psi - C)$$

Testing for confirmation and not predictively will therefore occur if $C > \Phi$. Together with the condition for testing for confirmation rather than not testing at all, we have $\Psi > \Phi$.

□

This model delivers the implication that testing is increasing in p directly through the wishful thinking. As in a neoclassical model, individuals want to test when they are most uncertain (when they believe the probability of the disease is 50%). Because of wishful thinking, this occurs at an objective probability higher than 50%, namely at p^*.

The other parts of Appendix Proposition 4 are more similar to the rational model. Individuals are prevented from testing here only by the real cost of testing, and therefore this cost must be high in order to prevent testing. In particular, for the same parameter values preventing testing requires a higher cost of testing than in the optimal expectations model. The required cost is not as large as in the neoclassical model, however, since there the value of testing was $p\Phi$, and here it is $\pi\Phi$ and, as shown in Appendix Proposition 3, $\pi < p$.

With regards to confirmatory testing, this model has an issue similar to the rational model. Generating confirmatory testing requires that the gain to confirmatory testing be larger than the gain to taking the correct actions up to that point. We have argued this is not the case, so effectively this model cannot rationalize confirmatory testing behavior.

Summary

In many ways this model is observationally similar to the optimal expectations case, given what we can see in our data. Both models predict biased beliefs, skewed actions and testing increasing in risk. In terms of the predictions on beliefs, this model may actually be a better fit than the optimal expectations, since it accommodates the fact that reported beliefs are increasing in actual risk, even if slowly, rather than suggesting all biased beliefs are $\pi = 0$.

The two models have different psychological micro-foundations, but testing between these is beyond the scope of our data. More specifically, the underpinnings of the optimal expectations model suggest that individuals "know" their true p – it's effectively a two-self model in which, if pushed, individuals could in principle access their true risk, even if they are adopting biased beliefs. The wishful thinking model has only the biased beliefs. Through wishful thinking individuals literally forget their true risk. Although this does differentiate the models, it's not something we can test in our data.

The most significant difference between this and the optimal expectations model comes in the testing. Like the neoclassical model, there is nothing limiting testing here other than real testing costs. This means that both (a) higher real testing costs would be required to rationalize low testing rates and (b) this model struggles to explain confirmatory testing. For this reason, we favor the optimal expectations setting, although the similarities between the models make this a plausible alternative. Richer data, with more information on psychological processing of these decisions, might help differentiate the two.

Information-Averse Preferences

Setup

We retain the basic setup from Section 4.1. The state is binary ($s \in \{0, 1\}$) and individuals have some exogenously given $p = E(s)$. Consumption utility at time 2 is given by:
$U_2 = E[u(a, s)|p] = pu(a, 1) + (1 - p)u(a, 0)$.

Individuals also experience some anticipatory utility at time 1. Individuals have information-averse preferences over anticipatory utility (Kreps and Porteus, 1978). We adopt an extremely simple formation for time 1 utility (or "anxiety"), drawn from Caplin and Leahy (2004): $U_1 = c(1 - p) - b\left(p - \frac{1}{2}\right)^2$, with $c, b > 0$. The $c(1 - p)$ term indicates an individual with these preferences would prefer a low p, since it indicates they are unlikely to be sick. However, this individual also (all else equal) has a preference for uncertainty: as p moves away from $\frac{1}{2}$ in either direction, utility falls. We can interpret c as a general disutility of being sick. We interpret b as the degree to which the individual values uncertainty.[2]

Since action choices *do not* enter time 1 utility, a^* is determined based only on consumption utility at time 2, so $a^*(p) = argmax_a[pu(a,1) + (1-p)u(a,0)]$. Total two period utility is given by the equation below.

$$U(a,s,p) = c(1-p) - b\left(p - \frac{1}{2}\right)^2 + pu(a^*(p),1) + (1-p)u(a^*(p),0)$$

We define the parameter values as in Section 4.

$$\begin{aligned} u(0,1) &= -\Omega \\ u(1,1) &= 0 \\ u(1,0) &= 1 - \Phi \\ u(0,0) &= 1 \end{aligned}$$

Results

We summarize the results in two propositions. The first describes beliefs and actions. The second describes the relationship between testing and risk. We discuss confirmatory testing at the end.

Appendix Proposition 5. Beliefs and Actions

1. *Self-reported beliefs are accurate* ($\pi = p$).
2. *Action* $a = 0$ *is taken as long as* $p < p^*$ *where* $p^* = \frac{\Phi}{\Phi+\Omega}$. *Note* $p^* > .5$ *iff* $\Phi > \Omega$.

Proof. 1. This follows directly from the fact that the model does not accommodate the possibility of beliefs other than p.

2. Note that the action utility enters only at time 2. The individual will choose $a = 0$ if $(1-p)\Phi \geq p\Omega$. Rearranging, we find that the individual will choose $a = 0$ up to a value of $p^* = \frac{\Phi}{\Phi+\Omega}$. Note this is the same as the neoclassical case.

□

This proposition notes first that this model cannot accommodate biased beliefs. Second, it demonstrates that generating skewed actions in this model requires the same type of asymmetry as in the neoclassical case, summarized in Section 5 in the paper.

Appendix Proposition 6. Testing and Risk *Define* $p^* = \frac{\Phi}{\Phi+\Omega}$, *as above. Individuals will choose to test if one of the following conditions holds:*

1. $p \leq p^*$ *and* $p > \frac{b-\Omega}{b}$
2. $p > p^*$ *and* $p < \frac{\Phi}{b}$

Proof. 1. When $p \leq p^*$ the individual is taking action $a = 0$. Given this, $V_{test} = p\Omega + b(p^2 - p)$. This is positive for values of $p > \frac{b-\Omega}{b}$.

[2] In principle, this taste for uncertainty could imply people actually prefer $p = .5$ to $p = 0$, which seems unlikely; however, with high enough c this will not be the case.

2. When $p > p^*$ the individual is taking action $a = 1$. Given this, $V_{test} = (1-p)(\Phi - bp)$. This is positive when $\Phi - bp > 0$.

□

For values of p below p^*, the utility from testing has two components. First, over this range of p, individuals take the action $a = 0$. As p increases, it becomes more and more likely this is wrong and the value to learning the true state and taking the correct action increases. This component of the testing utility is always positive. Pushing against this is the anxiety associated with testing which is the difference in anticipation utility if tested versus untested: $-\frac{b}{4} + b\left(p - \frac{1}{2}\right)^2$. This difference is the greatest (most negative) when $p = .5$. If b is large enough, this anxiety may be large enough to outweigh the positive benefits of testing and individuals may prefer not to test.

For small values of b, individuals with $p < p^*$ will always want to test. The model requires values of $b > \Omega$ in order to generate any testing avoidance over this range. Given that $b > \Omega$, testing value will be increasing with p up to p^*, which could produce the result that testing is increasing in risk; to have testing value be highest somewhere above 0.5 it must be the case that $b > 2\Omega$.

For value of $p > p^*$, the value of testing is decreasing in p. If $\Phi > b$ then the actor will want to test for any value of $p > p^*$. If $b > \Phi$ they may test only for values of p closer to p^*.

It is worth noting that, while this model does include the assumption that people prefer to have a lower p (in the form of the $c(1-p)$ term), this does not enter the decision-making about testing. Because this is linear in p, from the standpoint of someone deciding whether or not to test, individuals experience the same value for this (in expectation) whether or not they choose testing.

Confirmatory testing fits easily in this model. Individuals avoid testing here not due to testing costs but due to preferences for information avoidance. Once those concerns are removed (as they would be once the individual was sure of their status) they would be willing to test if Ψ (the value of confirmation) is greater than the cost of testing. This relationship has no bearing on the desire to avoid testing in the anticipation period.

Summary

Under some conditions, this model can fit the data. However, it shares two problems with the neoclassical model. First, it does not allow for overly-optimistic beliefs. To the extent this is a fact we would like to fit, the model fails there. In addition, explaining the skewed actions and testing increasing in risk requires some asymmetry in the loss to the wrong action in the two states (just as in the neoclassical model). As we discuss in Section 5, there is not a lot of empirical support for that claim. The latter result on testing and risk also requires a particular set of parameter values for b relative to Ω, an assumption which is plausible but difficult to test. Where this model clearly preforms better than the neoclassical model is on the testing costs and testing confirmation: it is able to explain limited testing without resorting to the extremely high testing costs and it can accommodate the confirmatory testing easily.