

Wishful Thinking*

Andrew Caplin[†] and John Leahy[‡]

June 2016

Abstract

We model agents who get utility from their beliefs and therefore interpret information optimistically. While subjectively Bayesian, they exhibit several biases observed in psychological studies such as optimism, confirmation bias, polarization, and the endowment effect.

1 Introduction

Expectation formation is central to many economic questions. Workers must form expectations regarding retirement. Investors must form expectations of risk and return. Price setters must form expectations of competitor's prices. While expectations are central, we do not fully understand how expectations are formed. The typical approach is to assume that expectations are model consistent. Agents understand the world in which they live and form expectations rationally. There is a lot of psychological evidence, however, that agents are poor information aggregators. It is therefore of interest to develop models of belief formation that go beyond the assumption of rational expectations.

In this paper, we proceed in the spirit of Becker and model beliefs as a choice. Our agents are subjective Bayesians. They accurately combine subjective signals with priors to form

*We thank Anmol, Bhandari, Kaitlin Raimi, Matthew Shapiro, Linda Tesar, and Jaume Ventura for helpful discussions.

[†]Center for Experimental Social Science and Department of Economics, New York University. Email: andrew.caplin@nyu.edu

[‡]Department of Economics and Gerald R Ford School of Public Policy, University of Michigan and NBER. Email: jvleahy@umich.edu

subjective posteriors. Their subjective reality is rational. Their problem is that they see the world through rose-colored glasses. They interpret signals optimistically to maximize their expected utility subject to a cost, which might be interpreted as a cost of fooling themselves. Hence their subjective interpretation of the signals differs from the objective counterpart. While subjectively Bayesian, our wishful thinkers will appear non-Bayesian to an objective observer.

Any model of belief choice must specify the costs and the benefits of distorting beliefs. In the standard economic model, there is no benefit to believing anything other than the truth. Agents get utility from outcomes, and probabilities serve only to weight these outcomes. Getting the probabilities wrong muddles one's view of the payoffs to an action and leads to mistakes. In the standard model, there are strong incentives to have accurate beliefs, since accurate beliefs lead to accurate decisions.

To model the benefits of belief choice we follow Jevons (1905), Loewenstein (1987) and Caplin and Leahy (2001) and assume that some portion of utility depends on the anticipation of future outcomes. Anxiety, fear, hopefulness and suspense are all ways in which beliefs about the future affect wellbeing today. The American Psychiatric Association defines anxiety as “apprehension, tension, or uneasiness that stems from the anticipation of danger.” The dependence of current wellbeing on beliefs about the future creates an incentive to believe that “good” outcomes are more likely than “bad” outcomes. Wishful thinking involves choosing to believe that the truth is what one would like the truth to be.

Without constraints a theory of belief choice would lack content. One could believe anything one liked. Our model of the constraints rests on the idea that there are often lots of possible beliefs that are consistent with experience. We limit our agents to “plausible” beliefs, by which we mean beliefs that are not obviously contradicted by the available evidence. We follow Hansen and Sargent (2008) and impose a cost to beliefs that are too far away from the truth, where we associate the truth with the beliefs that an objective observer would hold. This cost is related to the Kullback-Leibler divergence from the agents subjective beliefs to the objective probability distribution. The Kullback-Leibler divergence measures the likelihood that the subjective beliefs would be rejected in favor of the objective ones. model given current beliefs. The more unlikely the truth, the more costly the beliefs. We consider two formulations of this cost. In the first, which occupies much of the body of the paper, we place the cost on the interpretation of the flow of new information. This corresponds to the idea that agents “see the world through rose colored glasses.” In this formulation, the agent interprets signals sequentially. Later we also consider a formulation in which the agent evaluates all information at once. Each formulation has its advantages

and disadvantages, which we discuss in detail.

The outcome is a model of belief choice. Optimal beliefs tend to twist the probabilities in the direction of events with high utility. The upper bound on probabilities limits wishful thinking about very likely events. Events can only be so likely. While unlikely events with high payoffs receive more weight, wishful thinking is not magical thinking. Low probability events remain low probability and zero probability events remain zero probability. Wishful thinking is strongest when payoff differences are large and outcomes are uncertain. Such situations might plausibly include choices that are made infrequently so that the agent lacks experience with them such as saving for retirement. They might also include situations in which the options are difficult to value such as the valuation of real estate where every house is in some sense a unique asset and few houses trade at any given time. They might include any situation in which there are multiple theories on the table and very little evidence to distinguish between them as is often the case with asset bubbles.¹

While our wishful thinkers are subjective Bayesians, they may exhibit a number of deviations from rationality often observed in the experimental literature. Our wishful thinkers will tend to be overconfident. To the extent that more accurate decisions yield greater payoffs, they will overestimate the probability that they have made the correct decision and they will underestimate confidence intervals. Our wishful thinkers will tend to interpret information in ways that accord with their priors, a phenomenon known as confirmation bias. To the extent that their priors are the result of wishful thinking, both their priors and their interpretation of information will be twisted in similar ways by their desire to believe that the state is good. Two wishful thinkers with very different payoffs may interpret the same information in fundamentally different ways. Each may interpret the information in a way that is favorable to them. It is even possible that each becomes more confident that their view of the world is correct, a phenomenon known as polarization.

While wishful thinkers may appear non-Bayesian to an objective observer, their world view is internally coherent. This allows us to apply the standard decision theory tool kit. They maximize the present value of subjective expected utility. Conditional on their interpretation of information, they filter their subjective signals correctly.

Our premise is that people shade their beliefs in ways that make their choices look better. There is evidence that supports this assertion. In a classic study of cognitive dissonance, Knox and Inkster (1968) interviewed bettors at a race track and found that bettors placed

¹Reinhart and Rogoff title their study of credit booms "This Time is Different", emphasizing that there are multiple interpretations of the evidence and a tendency to gravitate towards the optimistic ones.

higher odds on their preferred horse when interviewed after placing their bets than bettors did when interviewed while waiting to place their bets. Knox and Inkster attribute this phenomenon to a desire to reduce post-decision dissonance, that is a desire to match one's world view with ones decisions. Mijovic-Prelec and Prelec (2010) perform a similar analysis in a more controlled setting. They had subjects make incentivized predictions before and after being given stakes in the outcomes, and found that there was a tendency for subjects to reverse their predictions when the state that they had predicted to be less likely turned out to be the high payoff state. Bastardi, Uhlmann, and Ross (2011) conduct a standard test of confirmation bias, but instead of focusing on the relationship between one's prior beliefs and their interpretation of evidence, they instead focus on the one's prior decisions. They consider a population of parents who profess to believe that home care is superior to day care for their children. Some, however, have chosen home care. Others, because they have jobs, have chosen day care. They present this population with two fictional studies, one of which claims home care is superior and one which claims day care is superior. The parents who had placed their children in day care rated the study supporting day care as superior, whereas the parents who cared for their children at home did the opposite. Some of the parents who had placed their children in day care changed their beliefs and professed day care was no worse than home care.

Section 2 discusses related literature. Section 3 presents a simple model of belief choice and discusses its implications for belief choice and action choice. We show that wishful thinking naturally leads to optimism, confirmation bias, polarization, and situations in which the order of information matters. Section 4 discusses a number of issues, including alternative modeling choices and the relationship to the literature on robust control. Section 5 concludes.

2 Related Literature

We contribute to three literatures. The first is the literature on belief choice which is surveyed by Benabou and Tirole (2016). They divide the literature into two classes depending on the motivation for distorting one's beliefs. In one class, beliefs enter directly into utility. In the other beliefs are instrumental in motivating desirable actions or achieving desirable goals. For example, beliefs may aid in overcoming self-control problems, signaling one's type, or fostering commitment. Our paper fits into the first class. Akerlof and Dickens (1982) is a prominent early example. They present the example of an agent considering a job in a hazardous industry. Upon accepting the job the agent may choose to understate the probability of an accident in their industry. This is desirable since it reduces fear,

and fear reduces utility. The cost of distorting beliefs is that mistaken beliefs may lead to suboptimal decisions in subsequent periods. For example, the agent may choose to forgo safety equipment if it became available. Akerlof and Dickens do not model the cost of distorting beliefs. Rather they constrain the subjective probability of an accident to be less than the objective probability. This leads to bang-bang solutions depending whether the cost or the benefit is greater. They assume that the agent evaluate these options with the true probabilities when considering belief choice. Brunnermeier and Parker (2005) is another closely related paper. They assume that an agent makes a once and for all belief choice in period zero and then given this prior behaves as a Bayesian in all subsequent periods. They model belief choice as balancing the gain to anticipating a more positive future and against the cost of suboptimal decisions. Like Akerlof and Dickens the agent evaluates belief choice using the true probability distribution. We discuss the relationship between our model and Brunnermeier and Parker’s in detail in Section 4 below.

A second literature is the literature on anticipatory feelings. Jevons (1905), Loewenstein (1987) and Caplin and Leahy (2001,2004) all suppose that current happiness depends in some way on future outcomes. Jevons believed that agents acted only to maximize current happiness. Intertemporal optimization, in this view, balanced the happiness from actions today against the current happiness arising from the anticipation of future actions. Loewenstein builds a model to explain why an agent might wish to bring forward an unpleasant experience to shorten the period of dread, or to postpone a pleasant experience in order to savor the anticipation. Caplin and Leahy model emotional responses to future risks such as anxiety, suspense, hope, and fear. Brunnermeier and Parker build on the utility function in Caplin and Leahy.

Our paper also contributes to the recent explosion of work that deviates from rational expectations. A partial and incomplete selection include the following. Hansen and Sargent (2008) consider robust expectations that incorporate a fear of model mis-specification. Fuster, Laibson, and Mendel (2010) propose what they call “natural expectations” which involve a weighted average of rational expectations and the prediction of a simple linear forecasting model. Gabaix (2014) considers a “sparsity-based” model in which agents place greater weight on variables that are of greater importance. Bordalo, Gennaioli, and Schleifer (2018) consider what they call “diagnostic expectations”. These are based on what psychologists call the representative heuristic and involve an overweighting of outcomes that are becoming more likely.

3 A Simple Model of Belief Choice

The essential elements of the theory are: (1) a decision whose outcome is unknown; (2) objective probabilities of the outcome; (3) utility from beliefs regarding the outcome; and (4) a cost to choosing beliefs that differ from the objective probabilities. We discuss these elements in turn.

There are two periods. In the first period, an agent chooses an action a from a finite set of potential choices A . In the second period, nature selects a state ω from a finite set of potential states of the world Ω . The agent begins the first period with a prior $\mu \in \Delta(\Omega)$ over second period states. The agent also observes a signal $s \in S$ in the first period. The objective probability of observing s when the state is ω , is $\bar{p}(s|\omega)$. The agent chooses the action a after observing the signal s . Unless otherwise noted we assume that $\bar{p}(s|\omega) > 0$ for all ω so that the signal is not perfectly informative.

We follow Jevons (1905) and assume that the agent maximizes current utility and that current utility incorporates the anticipation of future outcomes. This is the easiest way to incorporate utility from anticipation. Let $u(\omega, a)$ denote the utility that the agent receives from the anticipation of action a in state ω . The agent chooses the action a that maximizes the subjective expectation of $u(\omega, a)$.²

In addition to the choice a , the agent can choose how to interpret the signal s by choosing a subjective interpretation of the signal $p(s|\omega)$ that differs from the objective probability $\bar{p}(s|\omega)$. The benefit of choosing $p(s|\omega)$ is that it may increase the agent's subjective expected utility. The cost depends on how $p(s|\omega)$ differs from $\bar{p}(s|\omega)$. Rather than model the specific technology by which beliefs are distorted, for example by selective memory, selective attention, or self-signalling, we hypothesize that the costs of belief distortion are increasing in the size of the distortion.³ We first invert $p(s|\omega)$ and $\bar{p}(s|\omega)$, so that they become probabilities of ω given s . Let $\bar{p}(\omega) \equiv \bar{p}(s|\omega)/\bar{p}(s)$ and $p(\omega) \equiv p(s|\omega)/p(s)$. We then define the cost of choosing $p(\omega)$ as the Kullback-Leibler divergence from $\bar{p}(\omega)$ to $p(\omega)$. The information cost is

$$\frac{1}{\theta} \sum_{\omega} p(\omega) \ln \frac{p(\omega)}{\bar{p}(\omega)} \quad (1)$$

²The agent receives utility from the action a both through prior anticipation and eventual experience. Only the former influences choice. In dynamic models with multiple periods choice is time consistent if anticipatory utility mirrors experienced utility and the agent discounts the anticipation of future utility exponentially. Optimal policy involves additional complications. See Caplin and Leahy (2006).

³See Bénébou and Tirole (2016) for a discussion of theories of selective memory, selective attention and self-signalling.

This cost (1) is the expected likelihood ratio under the subjective measure p . It measures the ability of the agent to discriminate between p and \bar{p} given that the agent believes that the signal is p . The idea is that it is easier to choose a subjective belief that is not wildly contradicted by experience. θ is a parameter that captures the ease with which the agent can manipulate their beliefs. The larger is θ the greater the amount of evidence the agent would need before they reject their chosen beliefs in favor of the objective ones. We also do not take a stand on whether the process of twisting information is conscious or subconscious. Note that we place the cost on the interpretation of the signal rather than the posterior as do Hansen and Sargent. There are advantages to each approach which we discuss in Section 4 below.

Summarizing the above, the agent's maximization problem becomes

$$V(\mu, s) = \max_{p \in \Delta(\Omega), a \in A} \sum_{\omega} \frac{p(\omega)\mu(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} u(a, \omega) - \frac{1}{\theta} \sum_{\omega} p(\omega) \ln \frac{p(\omega)}{\bar{p}(\omega)}. \quad (2)$$

The state variables are the prior μ and the signal s . Given the interpretation of the signal, Bayes rule implies that the subjective posterior of state ω is

$$\gamma(\omega) \equiv \frac{\bar{p}(s|\omega)\mu(\omega)}{\sum_{\omega'} \bar{p}(s|\omega')\mu(\omega')} = \frac{p(\omega)\mu(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')}.$$

The first term is therefore the subjective expected utility of choice a . The second term is the cost of the believing p . Implicit in the maximization problem (2) is the assumption that the agent understands that their interpretation of the signal s will depend on the choice a . We consider the implications of assuming that the agent is naive in Section 4.

3.1 Implications for belief choice

Given that the agent understands the interaction between belief choice and action choice, it does not matter whether the agent chooses beliefs and then actions or actions and then beliefs. We will therefore fix the action and focus, for the time being, on belief choice. Given that the action a is fixed, we write $u(\omega)$ for $u(a, \omega)$. The first order condition for $p(\omega)$ implies

$$p(\omega) = \frac{\bar{p}(\omega) \exp \left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)} \right]}{\sum_{\omega'} \bar{p}(\omega') \exp \left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega')} \right]} \quad (3)$$

where $E_\gamma u(\omega)$ is the expectation of $u(\omega)$ with respect to the subjective posterior γ , and $\frac{\partial E_\gamma u}{\partial p(\omega)}$ is the partial derivative of this expectation with respect to $p(\omega)$. The derivation of (3) is in the appendix.

According to (3), the agent distorts their interpretation of the signal. They tend to increase the probability of states when increasing that probability increases expected utility and they reduce the probability of states when reducing that probability increases expected utility. The derivative $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}$, can be written as

$$\frac{\mu(\omega)}{E_p \mu(\omega)} [u(\omega) - E_\gamma u(\omega)].$$

According to the term in brackets, the agent tends to raise the probability of states with above average utility (according to the posterior γ). This is the essence of wishful thinking. The agent believes to be true what they would like to be true. According to the ratio in front of the term in brackets, the agent tends to distort beliefs more (in absolute value) if the prior probability of the state is high. Given the cost of distorting beliefs, it does not make sense to waste effort distorting unlikely events. $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}$ in (3) is multiplied by θ . The larger is θ , the easier it is for the agent to manipulate their interpretation of the signal.

Note that since p affects γ , it enters both sides of (3). These first-order conditions therefore implicitly define p . The next proposition shows that a solution always exists, so (3) is not vacuous. All proofs are in the appendix.

Proposition 1 Given $\mu, \bar{p} \in \text{int}\Delta(\Omega)$, there exists a $p \in \text{int}\Delta(\Omega)$ that satisfies (3) for all $\omega \in \Omega$.

It is possible that there are multiple solutions to (3). This can happen when $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}$ is increasing in $p(\omega)$, then an increase in $p(\omega)$ raises both the gain in subjective utility and the cost of belief distortion. In such cases, the agent chooses the solution associated with the highest $V(\mu, s)$. In most simulations, this solution has been associated with the most optimistic beliefs.

We now consider two special cases which illustrate the how the model works.

3.1.1 Example with uniform priors

Suppose that the prior is uniform $\mu(\omega) = 1/N$ where $N = |\Omega|$. In this case it is easy to show that $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} = u(\omega)$, so that the marginal increase in expected utility from increasing

the probability of state ω is simply the utility in that state. The first-order condition (3) simplifies dramatically

$$p(\omega) = \frac{\bar{p}(\omega) \exp[\theta u(\omega)]}{\sum_{\omega'} \bar{p}(\omega') \exp[\theta u(\omega')]} \quad (4)$$

Note that (4) pins down $p(\omega)$ uniquely. The optimal choice of p distorts the true probability distribution \bar{p} in the direction of the more desirable states.

3.1.2 Example with two states

Suppose that there are two states ω_H and ω_L with $u_H \equiv u(\omega_H) > u(\omega_L) \equiv u_L$ so that ω_H is the good state and ω_L is the bad state. Let p_H and p_L denote the optimal subjective beliefs of the respective states and define \bar{p}_H , \bar{p}_L , μ_H and μ_L similarly. With these definitions (3) can be written as,

$$p_H = \frac{\bar{p}_H}{\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}} \quad (5)$$

In this case, p_H still appears on both sides of the equation, but the assumption of two states eliminates much of the interaction between states and allows us to state the comparative static results cleanly. We collect these in the next proposition.

Proposition 2 With two states:

1. $p_H > \bar{p}_H$
2. p_H is strictly increasing in $u_H - u_L$.
3. p_H is strictly increasing in θ .
4. p_H is strictly increasing in \bar{p}_H .
5. p_H is strictly increasing in μ_H if $\mu_H < 1 - p_H$ and decreasing in μ_H if $\mu_H > 1 - p_H$.

Point (1) is the essence of wishful thinking: the subjective interpretation of the signal is more optimistic than the objective interpretation, which implies that the subjective posterior will also be more optimistic. Point (2) states that the extent of wishful thinking is increasing in the relative payoff of the desirable state. Point (3) states that wishful thinking is decreasing in cost parameter $1/\theta$. Point (4) reflects the effect of the objective probabilities on the subjective probabilities. Finally, point (5) reflects the sensitivity of the posterior with respect to the signal.

Two states also allows us to sharpen our characterization of when (3) has a unique solution.

Proposition 3 With two states:

1. If $\mu_H > \mu_L$, then there is only one solution to (5).
2. If $\theta(u_H - u_L) < 16$ then there is only one solution to (5).

Point 1, $\mu_H > \mu_L$, is sufficient for the right-side of (5) to be decreasing in p_H . Point 2 is sufficient for the derivative of the right-side of (5) to be less than one. In general, multiple solutions are more likely if θ or $u_H - u_L$ are high and μ_H is low. All of these increase the response of $E_\gamma u(\omega)$ to p_H .

Figure 1 graphs $p(\omega)$ as a function of $\bar{p}(\omega)$ for an example with two states and a uniform prior. The prior probability of the high utility state ω_H is on the horizontal axis. The solid line represents the chosen probability as measured on the vertical axis. The gap between the solid line and the dashed 45 degree line represents the extent of wishful thinking, $\left(\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}\right)^{-1}$. Since the state is good, this gap is everywhere positive. The constraint that probabilities are less than one, limits the amount of wishful thinking when the good state is very likely. It is hard to be over-optimistic about a near certain event. For example, Germans may be optimistic about their team's chances to win the FIFA world cup, but this does not necessarily reflect wishful thinking. Similarly, it is hard to be too optimistic about very uncertain events. Wishful thinking in this case is not magical thinking. If \bar{p}_H is zero, then p_H and so will be γ_H . Once the US team has been eliminated from the world cup, it is difficult to for Americans to fantasize about their winning the tournament.

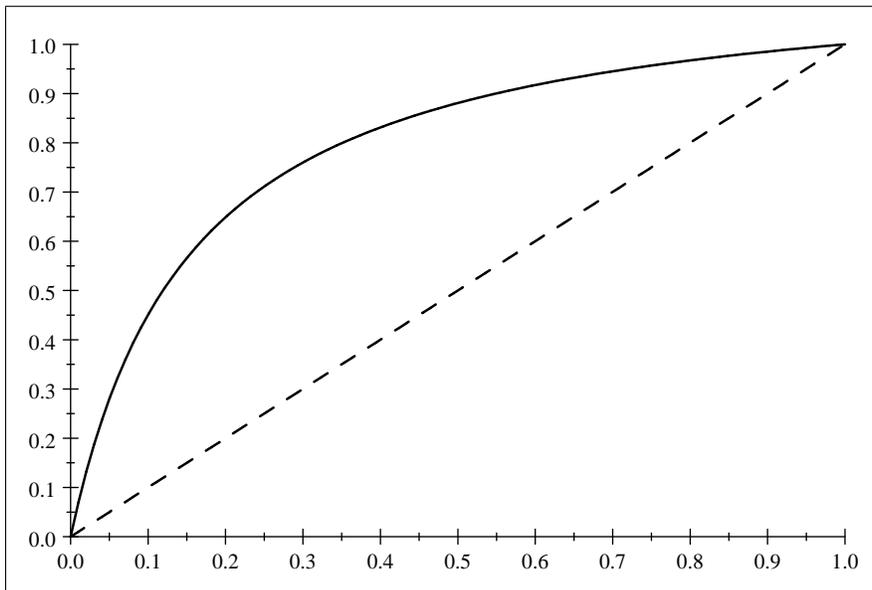


Figure 1: The figure depicts p_H as a function of \bar{p}_H for $\mu_L = \mu_H = .5$ and $\theta(u_H - u_L) = 2$.

With uniform priors, the posterior is equal to the interpretation of the signal: $\gamma_H = p_H$. In this case, Figure 1 also tells us how the subjective posteriors differ from the objective posteriors.

Wishful thinking is strongest when the signal is ambiguous. The situations in which we are likely to see wishful thinking are situations in which it is difficult to know the value of an option, such as the purchase of a house, or the agents have little experience, such as retirement, or there are multiple theories on the table, such as when there are trends in the data.

It is easy to show that in this example with uniform priors, the difference between p_H and \bar{p}_H peaks at $\bar{p}_H < \frac{1}{2}$, so that wishful thinking is strongest when the good state ω_H is slightly less likely than the bad state. There is a range of \bar{p}_H for which $\bar{p}_H < \frac{1}{2}$ and $p_H > \frac{1}{2}$. This can help explain why people gamble. Note that the house can exploit this behavior to a certain extent but not too much. The odds can favor the house but $\mu(\omega^G)$ must stay sufficiently positive that agents can believe $\gamma(\omega^G) > \frac{1}{2}$.

Figure 2 illustrates the effect of the prior. The black line shows p_H as a function of \bar{p}_H when the prior is uniform and the parameters are as in Figure 1. The green line shows the effect of increasing μ_H to .75. As stated in Proposition 2, the increase in μ_H increases p_H for low values of \bar{p}_H and reduces p_H for high values of \bar{p}_H . The red line shows that reducing μ_H to .25 has the opposite effect.

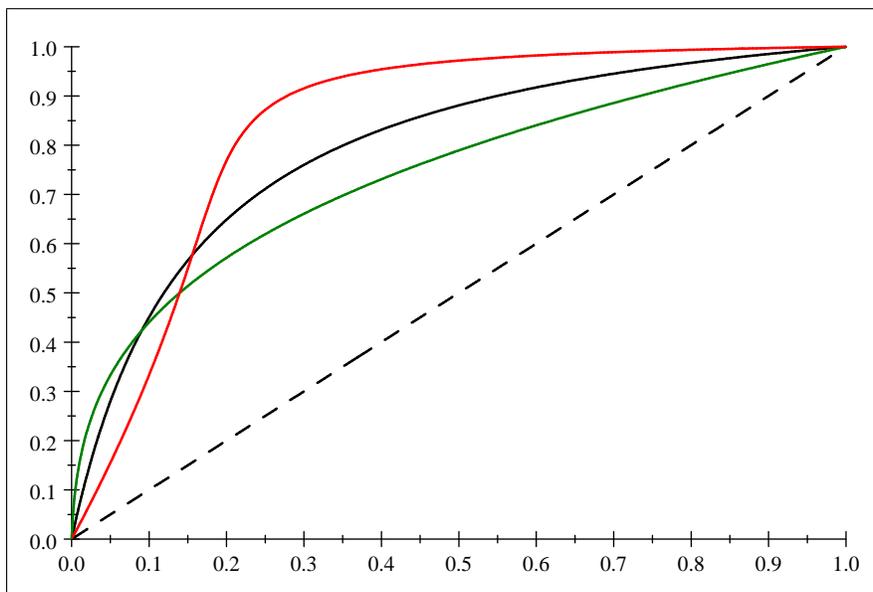


Figure 2: The effect of the prior on the interpretation of the signal.

This effect on the prior carries over to the posterior. Figure 3 shows the resulting posteriors as function of \bar{p}_H . The green line in Figure 3 is slightly above the green line in Figure 2 because the higher prior leads to a higher posterior. The red line in Figure 3 is lower than the red line in Figure 2 for similar reasons. Overall, the posteriors react in much the same way as the interpretation of the signal.

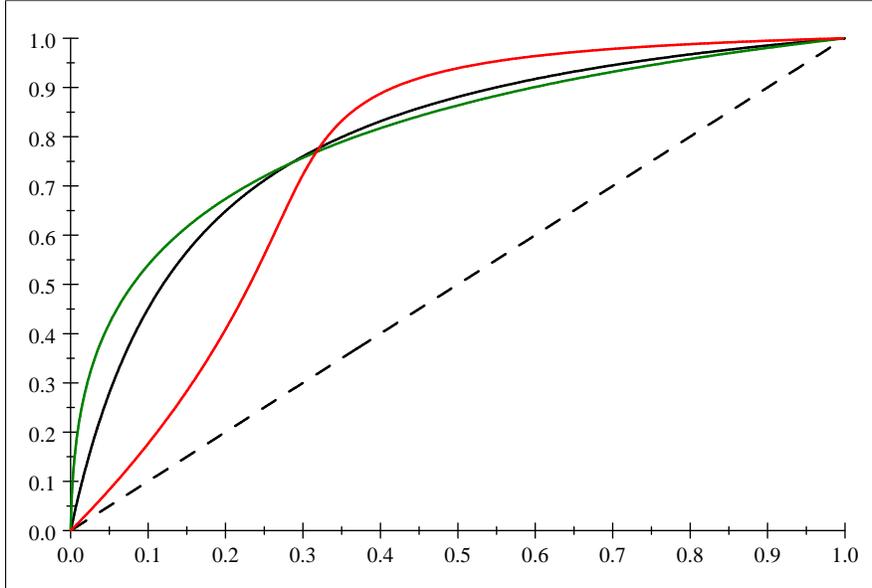


Figure 3: γ_H as a function of \bar{p}_H .

3.2 Action Choice

We can calculate the value of the action a under the optimal beliefs. Substituting (3) into $V(\mu, s)$ for a given action choice a ,

$$V(\mu, s) = \max_a E_{\gamma(a)} u(a, \omega) + \frac{1}{\theta} \left(\ln \sum_{\omega} \bar{p}(\omega) e^{\theta \frac{\partial E_{\gamma(a)} u(\omega)}{\partial p(\omega)}} \right)$$

In the case with uniform priors this simplifies to

$$V(\mu, s) = \max_a \frac{1}{\theta} \left(\ln \sum_{\omega'} \bar{p}(\omega') e^{\theta u(\omega')} \right)$$

This has the form of Epstein-Zin (1989) preferences, $f^{-1}(E\{f(u(\omega))\})$ where $f(x) = \exp(x)$.

Given that $\exp(x)$ is convex, the agent has a preference for late resolution to uncertainty. This is not surprising as it is uncertainty that allows the agent to engage in wishful thinking.

3.3 Implications

Our wishful thinkers are subjective Bayesians, but they will appear to be non-Bayesian to an objective observer. In this section, we argue that many apparent deviations from Bayes rule such as overconfidence, conformation bias, polarization, the endowment effect, and the effect of the order of information may be natural implications of wishful thinking.

3.3.1 Overconfidence

Debondt and Thaler (1995) write, “Perhaps the most robust finding in the psychology of judgement is that people are overconfident”. Experimental tests of overconfidence take several forms. In one type of experiment, an agent is asked to choose the correct answer from a set of potential answers and then asked their subjective probability of getting choosing the correct answer. In this case, overconfidence takes the form of optimism and arises when the subject’s subjective probability of being correct exceeds the observed frequency with which they in fact answer correctly. Another set of experiments asks for a numerical answer to a question and for a subjective confidence interval. In this case, overconfidence takes the form of excess precision and arises when the correct answer fails to lie in the subjective confidence interval as often as believed.

The spirit of these tests can be captured in a tracking problem in which the agent must guess the state after receiving a signal. Suppose that there are N states labeled ω_1 through ω_N equally spaced around a circle, so that the distance between ω_1 and ω_2 is equal to the distance between ω_1 and ω_N . Nature picks the true state $\hat{\omega} \in \Omega$, and the agent picks $a \in \Omega$. Let $\delta(a, \hat{\omega})$ denote the minimum distance (about the circle) between the true state and the choice and suppose that u depends only on δ : $u(a, \hat{\omega}) = u(\delta(a, \hat{\omega}))$. This payoff function captures both types of experiment. In the first case, there is a correct answer and a collection of incorrect answers. For example, $u(\delta) = 1$ if $\delta = 0$ and $u(\delta) = 0$ otherwise. In the second case, the loss is increasing in the δ . Suppose that the prior is uniform and that the signal has an objective distribution that is symmetric about the true state. The symmetry of the signal implies that $\bar{p}(\omega|\omega = s + x) = \bar{p}(\omega|\omega = s - x)$ for all x . Given the symmetry of the problem, the optimal choice is obvious: the agent simply reports the signal: $a = s$. The question that we focus on is what the agent chooses to believe.

Given the uniform prior, (3) becomes

$$p(\omega) = \frac{\bar{p}(\omega) \exp [\theta u(\delta(s, \omega))]}{\sum_{\omega'} \bar{p}(\omega') \exp [\theta u(\delta(s, \omega'))]}. \quad (6)$$

Without loss of generality label the signal state ω_0 . Label the states to the right of ω_0 (as we move about the circle): $\omega_1, \omega_2, \dots$, and label the states to the left $\omega_{-1}, \omega_{-2}, \dots$. If there are an odd number of states keep the number of states with positive and negative indices equal. If there are an even number of states, label the state furthest from ω_0 , $\omega_{|\Omega/2|}$. With this labeling, the absolute value of the index is equal to $\delta(s, \omega)$. Now consider the ratio

$$\frac{p(\omega_m)}{p(\omega_n)} = \frac{\bar{p}(\omega_m) \exp[\theta u(|m|)]}{\bar{p}(\omega_n) \exp[\theta u(|n|)]}$$

First, $\bar{p}(\omega_n) = \bar{p}(\omega_{-n})$ implies $p(\omega_n) = p(\omega_{-n})$, so that p inherits the symmetry of \bar{p} and δ . Second, since u is maximized at $\delta = 0$, $\frac{p(\omega_0)}{p(\omega_n)} > \frac{\bar{p}(\omega_0)}{\bar{p}(\omega_n)}$ for all $\omega_n \neq \omega_0$. It follows immediately that $p(\omega_0) > \bar{p}(\omega_0)$, so that the agent is over-optimistic that they have selected the correct state. Finally, if we consider any subset of states relatively close to ω_0 , $\bar{\Omega} = \{\omega_n | n < \bar{N}\}$, $\frac{p(\omega_n)}{p(\omega_m)} > \frac{\bar{p}(\omega_n)}{\bar{p}(\omega_m)}$ for any $\omega_n \in \bar{\Omega}$ and $\omega_m \notin \bar{\Omega}$. Hence the agent will be overconfident that the true state is in $\bar{\Omega}$. It follows that the agent will be over-confident in the sense that their subjective confidence intervals will be too tight.

3.3.2 A simple model of entry

To see the effects of over confidence consider a simple model of entry. A firm contemplates entry into a market. Upon entry the firm will produce a unit flow of a consumption good. There is no option of exit. In period one, the price of the consumption good is normalized to one. In period 2, the price will either rise to $q_H > 1$ or fall to $q_L < 1$ where it will remain forever. The firm receives a signal in period 1. Given the signal the objective probability that the price will be q_H is \bar{p}_H . The firm discounts future revenue by a factor β . There is an entry cost k . We assume that $k > 1 + \frac{\beta}{1-\beta}q_L$ so the firm does not wish to enter if the price falls.

Using the objective probabilities the firm would enter in period 1 if

$$1 + \frac{\beta}{1-\beta} [\bar{p}_H q_H + (1 - \bar{p}_H) q_L] - k > \beta \bar{p}_H \left[\frac{q_H}{1-\beta} - k \right]$$

The left-hand side is the expected value of immediate entry. The right-hand side is the expected value of waiting. This condition reduces to

$$1 > (1 - \beta)k + \beta(1 - \bar{p}_H) \left(k - \frac{q_L}{1 - \beta} \right)$$

The left-hand side is the lost revenue from waiting. The right-hand side is the gain from delaying the entry cost and the losses that can be avoided should the firm refrain from entering in the low state. It is immediate that entry is more likely the higher is \bar{p}_H and that q_H has no effect on the entry decision.

Now consider how the wishful thinker will behave. Suppose that their prior is uniform. There are two states ω^H is associated with the high price and ω^L with the low price. If the firm enters then the payoff in the high state is

$$1 + \frac{\beta}{1 - \beta} q_H - k$$

and in the low state is

$$1 + \frac{\beta}{1 - \beta} q_L - k$$

Their optimal belief conditional on entry is therefore

$$p_H^E = \frac{\bar{p}_H}{\bar{p}_H + (1 - \bar{p}_H) e^{\frac{\beta\theta}{1-\beta}(q_L - q_H)}}$$

Similarly we can calculate subjective beliefs conditional on waiting. In this case, the payoff in the high state is

$$\beta \left(\frac{1}{1 - \beta} q_H - k \right)$$

and in the low state is zero. Their optimal belief conditional on entry is therefore

$$p_H^W = \frac{\bar{p}_H}{\bar{p}_H + (1 - \bar{p}_H) e^{\beta\theta(k - \frac{1}{1-\beta} q_H)}} \in (\bar{p}_H, p_H^E)$$

p_H^W is greater than \bar{p}_H but less than p_H^E as the desire to believe in the high state is mitigated by the option to remain inactive should the low state occur.

The entry criterion becomes

$$1 + \frac{\beta}{1 - \beta} (p_H^E q_H + p_L^E q_L) - k > p_H^W \beta \left(\frac{1}{1 - \beta} q_H - k \right)$$

and rearranging

$$1 + (p_H^E - p_H^W) \frac{\beta}{1 - \beta} q_H > (1 - \beta)k + \beta(1 - p_H^W)k - \beta(1 - p_H^E) \frac{q_L}{1 - \beta}.$$

Wishful thinking has several effects on the entry condition. First, since $p_H^E > p_H^W$ there is an added gain to immediate entry. The left-hand side is greater than before. Second, since both p_H^E and p_H^W are less than \bar{p}_H the perceived gain to avoiding entry in the bad state is less. The right-hand side is less than before. It follows immediately that the firm is more likely to enter. Moreover, entry depends positively on the price in the good state, as this price affects the desirability of distorting beliefs.

3.3.3 Confirmation bias

Confirmation bias occurs when an agent interprets information in a way that confirms their priors. Wishful thinking occurs when an agent interprets information in a way that enhances their subjective utility. The connection between wishful thinking and confirmation bias rests on the observation that the agent's prior is itself the result of wishful thinking in the past and hence likely correlated with payoffs.

To illustrate confirmation bias, consider the example with two states, ω_1 and ω_2 . Consider two agents, one of whom receives high utility from state ω_1 and the other receives high utility in state ω_2 . According to Proposition 2, both will twist signals in the direction of their preferred state. Now suppose that the agents unexpectedly receive additional signals. Again applying Proposition 2, each agent will again twist the signals in the direction of their preferred state, which will also be the direction of their priors.

Most tests of confirmation bias take the priors as given and evaluate how an agent interprets additional information. They do not consider the agent's subjective utility. It is therefore difficult to know whether the interpretation is being influenced by the prior beliefs or whether both beliefs and the interpretation are being influenced by payoffs. A few studies attempt to disentangle the effects of beliefs and payoffs. Mijovic-Prelec and Prelec (2010) had subjects make incentivized predictions before and after being given stakes in the outcomes. There was a tendency for subjects to reverse their predictions when the state that they had predicted to be less likely turned out to be the high payoff state. Bastardi, Uhlmann, and Ross (2011) consider a population of parents with similar priors: all profess to believe that home care is superior to day care for their children. They differ, however, in their payoffs, as some have chosen home care for their children, while others have chosen day-care. They find that the interpretation of evidence aligns with the payoffs rather than the prior. The parents who had placed their children in day care rated the study supporting day care as superior, whereas the parents who cared for their children at home did the opposite. In both of these studies, the interpretation of information appears to be more responsive to payoffs

than priors. This does not imply that priors do not matter, but only that wishful thinking might be present as well.

3.3.4 Polarization

Polarization occurs when two agents with opposing beliefs see the same signal and each becomes more convinced that their view is the correct one. Wishful thinkers can exhibit polarization if they place different values on the states and the information that they receive is sufficiently ambiguous. Consider again a situation with two agents and two states. Suppose that agent i receives utility u_H in state ω_1 and u_L in state ω_2 and agent j receives utility u_H in state ω_2 and u_L in state ω_1 . In keeping with our discussion of confirmation bias, suppose that each has received some information in the past that they have interpreted optimistically, so that agent i has a prior that places weight $\mu_H > \frac{1}{2}$ on state ω_1 , and agent j places the same prior on state ω_2 . Each then sees the same signal s which has objective probability \bar{p} . Each interprets the signal according to (5). Agent i ends up with the posterior

$$\gamma_i(\omega_1) = \frac{\mu_H \bar{p}(\omega_1)}{\mu_H \bar{p}(\omega_1) + (1 - \mu_H) \bar{p}(\omega_2) e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}}$$

and agent j ends up with the posterior

$$\gamma_j(\omega_1) = \frac{(1 - \mu_H) \bar{p}(\omega_1)}{(1 - \mu_H) \bar{p}(\omega_1) + \mu_H \bar{p}(\omega_2) e^{\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}}$$

Polarization occurs if $\gamma_i(\omega_1) > \mu_H$ and $\gamma_j(\omega_1) < 1 - \mu_H$. so that both agents have observed the same signal and each has become more confident in thier assessment of the state. Now $\gamma_i(\omega_1) > \mu_H$ if $\frac{\bar{p}(\omega_2)}{\bar{p}(\omega_1)} e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}} < 1$. Similarly $\gamma_j(\omega_1) < 1 - \mu_H$ if $\frac{\bar{p}(\omega_2)}{\bar{p}(\omega_1)} e^{\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}} > 1$. Since $e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}} < 1$ and $e^{\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}} > 1$, it follows immediately that polarization is possible, and that polarization is more likely (at least in this example) when the objective odds $\frac{\bar{p}(\omega_2)}{\bar{p}(\omega_1)}$ are close to even and when $u_H - u_L$ is large. In other words, polarization tends to occur with the signal is uninformative and the desire to believe is large.

3.3.5 Order of information

Consider an agent who receives a signal, interprets it according to (3) and then receives a second signal, which they also interpret according to (3). Suppose that the two signals when

considered together are uninformative. Does the order of information matter?

It is clear that if the agent makes a decision between receiving the two signals, then the order of information will matter. The first signal will affect the agent's choice and that choice will influence the interpretation of both signals. Change the order of the signals and one would likely alter the choice and hence the interpretation.

Yet even if there is no choice the order of the signals will matter. Consider a wishful thinker how receives two signals in succession. Suppose that the second signal is unanticipated, since whether a signal is anticipated or unanticipated would not affect a Bayesian. For simplicity suppose that the prior is uniform and that there are two states

The posterior after the first signal is

$$\gamma_H^1 = \frac{\frac{1}{2}p_H^1}{\frac{1}{2}p_H^1 + \frac{1}{2}p_L^1} = p_H^1 = \frac{\bar{p}_H \exp(\theta u_H)}{\bar{p}_H \exp(\theta u_H) + \bar{p}_L \exp(\theta u_L)}$$

The posterior after the second signal is

$$\begin{aligned} \gamma_H^2 &= \frac{p_H^2 \gamma_H}{p_H^2 \gamma_H^1 + p_L^2 \gamma_L^1} \\ &= \frac{\frac{\bar{p}_L}{(p_H^2 \gamma_H^1 + p_L^2 \gamma_L^1)^2} \frac{\bar{p}_H \exp(\theta u_H)}{\bar{p}_H \exp(\theta u_H) + \bar{p}_L \exp(\theta u_L)}}{\frac{\bar{p}_L + \bar{p}_H e^{-\theta \gamma_H^1 \gamma_L^1 (u_H - u_L)}}{(p_H^2 \gamma_H^1 + p_L^2 \gamma_L^1)^2}} \\ &= \frac{\frac{\bar{p}_H \exp(\theta u_H)}{\bar{p}_H \exp(\theta u_H) + \bar{p}_L \exp(\theta u_L)}}{\frac{\bar{p}_H \exp(\theta u_H)}{\bar{p}_H \exp(\theta u_H) + \bar{p}_L \exp(\theta u_L)} + \frac{\bar{p}_L \exp(\theta u_L)}{\bar{p}_H \exp(\theta u_H) + \bar{p}_L \exp(\theta u_L)}} \end{aligned}$$

4 Discussion

4.1 Sophistication vs naïveté

We have chosen to model sophisticated agents that are aware of how choices affect their beliefs. When choosing an action in the decision problem (2), the agent foresees that their beliefs will change and takes this into consideration. Another possibility is that agents are naive. They may not consider or may not be aware of how their choices affect their beliefs.

Most of the phenomenon considered above would still be apparent if the agent were naive. Optimism, overconfidence and polarization only depended on the choice of beliefs, not the choice of actions. naïveté, however, gives rise to additional phenomenon such as the foot-in-the-door technique .

The “foot-in-the-door technique” involves getting a person to make a big decision by first

having them make a similar decision on a smaller scale. As an example consider a world with two states ω_H and ω_L . Suppose ω_H is more likely, $\mu_H > \mu_L$. Consider a gamble in which the agent gets x if the state is ω_H and $-x$ otherwise. Suppose that utility, $u(x)$, is increasing and concave. Then it is possible that the agent would choose the gamble for small x and avoid the gamble for larger x . It is also possible that upon choosing the gamble for small x , the agent's belief γ_H would rise enough that the agent would now be willing to take the larger gamble.

4.2 An alternative cost function

We have chosen to place the information cost on the interpretation of the signal. Agents interpret signals and then update their beliefs. The interpretation is potentially biased. The updating process is Bayesian. Agents affect their posterior beliefs by manipulating the flow of information.

An alternative approach would be assume that agents can manipulate the stock of information not only the flow. This alternative approach suggests placing the information cost on the posterior itself: world views that deviate from objective reality are costly. This approach is equivalent to assuming that the agent reevaluates all of their past signals simultaneously each period. Our approach above, in contrast, is to consider each signal sequentially.

A cost function that places the cost on the posterior would look like:

$$\sum_{\omega} \gamma(\omega) \ln \left(\frac{\gamma(\omega)}{\bar{\gamma}(\omega)} \right) \quad (7)$$

where $\bar{\gamma}$ is the objective interpretation of the state given the history of signals and $\gamma(\omega)$ is the subjective posterior. This is the cost function employed by Hansen and Sargent albeit to model pessimism rather than optimism. The agent's maximization problem becomes

$$V(\bar{\gamma}) = \max_{\gamma \in \Delta(\Omega), a \in A} \sum_{\omega} \gamma(\omega) u(a, \omega) - \sum_{\omega} \gamma(\omega) \ln \left(\frac{\gamma(\omega)}{\bar{\gamma}(\omega)} \right).$$

Here we have replaced the signal s and the prior μ with the objective posterior $\bar{\gamma}$ which aggregates all past information. The optimal choice of beliefs then implies:

$$\gamma(\omega) = \frac{\bar{\gamma}(\omega) \exp[\theta u(a, \omega)]}{\sum_{\omega'} \bar{\gamma}(\omega') \exp[\theta u(a, \omega')]} \quad (8)$$

and the value choosing action a becomes

$$V(a) = \frac{1}{\theta} \ln \left(\sum_{\omega} \bar{\gamma}(\omega) \exp [\theta u(\omega)] \right) \quad (9)$$

These equations are the same as those in Hansen and Sargent except for the sign of θ .

We will refer to our earlier approach with the cost function (1) as the sequential cost model and the approach in this section with the cost function (7) as the cumulative cost model.

There are advantages and disadvantages of each modelling approach. Placing the cost on posterior beliefs as in the cumulative cost model leads to an algebraically simpler solution which may prove useful in dynamic applications. The cumulative cost model can also explain many of the observed deviations from rational behavior that are explained by the sequential cost model. For example, the agents in the cumulative cost model are optimistic and overconfident.

The cumulative cost model can also explain some phenomenon that the sequential cost model cannot. For example, the cumulative cost model can explain the endowment effect: that a person values and object more highly when it is in their possession. The cumulative cost model explains the endowment effect through a revision in beliefs. Once in possession of the object, the agents payoffs change and this leads directly to a revision of the posteriors (8). Since the agent in the sequential cost model manipulates the flow of information and not the stock of information, the sequential cost model can only explain the endowment effect if the person who receives the object also receives a signal that they can manipulate.

There are also disadvantages to the cumulative cost model. It has more difficulty explaining confirmation bias and polarization. The reason is that subjective posteriors are closely tied to objective posteriors in (8). News that raises $\bar{\gamma}(\omega)$ will tend to raise $\gamma(\omega)$ for all agents.

There are other differences in the two approaches. Agents in the sequential cost model are subjective Bayesians and appear non-Bayesian to an objective observer. Agents in the cumulative cost model maximize (9). They therefore appear to be objective Bayesians with an Epstein-Zin utility function. If one attempts to elicit their subjective beliefs, however, these subjective beliefs will appear non-Bayesian.

Beliefs are more stable in the sequential cost model. They evolve with the flow of information. Beliefs can potentially change dramatically in the cumulative cost model. If an

agent chooses action a , and the payoff to action a changes, then the agent will alter their beliefs even if they have not received any new information.

4.3 Robustness or Wishful Thinking

Hansen and Sargent (2008) model agents as pessimistic. Their agents are concerned that their model of the economy is inaccurate, and seek to make sure that their decisions are robust to plausible alternatives. This leads to an optimization problem very similar to (??), but the cost (7) enters with the opposite sign and the agent first minimizes with respect to beliefs before maximizing with respect to actions. Not surprisingly, this leads to very different behavior. The agent distorts beliefs toward the low payoff states instead of the high payoff states, and behaves as if they have a preference for early resolution of uncertainty rather than late resolution of uncertainty.

Which model is a better model of human behavior is not an easy question to answer. Each model is supported by its own body of psychological evidence and each performs well on in certain domains and poorly in others. The psychological justification for robustness is that it is consistent with ambiguity aversion and generates a preference for late resolution of uncertainty which many find plausible. The economic justification for robustness is that a preference for robustness generates risk sensitive preferences which help to explain the behavior of asset prices, in particular the equity premium.

As discussed above, there is also psychological evidence that agents distort beliefs in the direction of payoffs, and the psychological evidence in favor of optimism is at least as strong as that in favor of ambiguity. And while robustness appears to help explain asset pricing behavior, there are many economic situations where are better explained by wishful thinking. Entrepreneurs, for example, appear optimistic. According to Daniel Kahneman, “A lot of progress in the world is driven by the delusional optimism of some people.” Cooper, Woo, and Dunkelburg (1988) find that two thirds of entrepreneurs believe that their firm will fare better than similar firms run by others. Hamilton (2000) finds that the median earnings of entrepreneurs is 35% less than would they would be predicted to earn in alternative jobs. Dropping out of Harvard to develop a social networking site as Mark Zuckerberg did would appear much more consistent with optimism than a preference for robustness.

Payday lending is another area that would appear more consistent with optimism than robustness. Payday loans typically accrue about 18% over a period of two weeks or an annualized value of over 7000%. Borrowers appear to be overoptimistic regarding their

ability to repay and end up rolling loans over multiple times. Borrowers tend to be optimistic regarding how many times they will roll over debt. Finally politicians rarely crow about the robustness of their policies, preferring instead to emphasize the optimistic outcomes.

It is a question for future research when people exhibit a preference for robustness and when people engage in wishful thinking.

4.4 Brunnermeier and Parker

The most closely related paper is Brunnermeier and Parker (2005). That paper like our paper presents a model of belief choice in which the benefit of belief choice is that beliefs enter directly into utility. Their agents like our agents are subjective Bayesians. Their model leads to many of the same phenomenon such as optimism and overconfidence.

There are, however, several differences. First, the two papers model the costs of belief choice in very different ways. Brunnermeier and Parker focus on how optimistic beliefs might lead to suboptimal decisions. In their model there is an initial period in which the agent chooses their prior. In subsequent periods, the agent observes the world, updates their information as would a Bayesian and makes decisions. The period-zero agent balances the utility gain from choosing an optimistic prior against the against the mistakes that result from this mistaken prior. They use the objective probabilities to weight outcomes. Our agents do not consider the costs that arise from mistaken beliefs. Instead we place the cost in how far their beliefs deviate from the objective evidence.

Second, while agents in both models are subjective Bayesians, they deviate from objective Bayesians in different ways. In Brunnermeier and Parker, agents have an incorrect prior, but their interpretation of evidence accords with objective reality. In our model, agents may or may not have an incorrect prior, it is there interpretation of signals that is overly optimistic.

Third, Brunnermeier and Parker model a once and for all choice of beliefs. All belief choice occurs in the initial period. Our agents twist each and every signal that they receive. In Brunnermeier and Parker beliefs affect choices. In our model, choices also mold beliefs.

5 Conclusion

We model an agent who get utility from their beliefs and therefore interprets information optimistically. While subjectively Bayesian, the agent exhibits several biases observed in

psychological studies such as optimism, confirmation bias, polarization, and the endowment effect.

References

- Akerlof, George, and William Dickens (1982), "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 72, 307-319.
- Bastardi, A.; Uhlmann, E. L.; Ross, L. (2011), "Wishful Thinking: Belief, Desire, and the Motivated Evaluation of Scientific Evidence," *Psychological Science*, 22, 731-732.
- Benabou, Roland, and Jean Tirole (2002), "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics* 117, 871-915.
- Benabou, Roland, and Jean Tirole (2016), "Mindful Economics: The Production, Consumption, and Value of Beliefs," *Journal of Economic Perspectives* 30, 141-164.
- Bernardo, Antonio, and Ivo Welch (2001), "On the Evolution of Overconfidence and Entrepreneurs," *Journal of Economics and Management Strategy*, 10, 301-330.
- Brownstein, Aaron, Stephen Read, and Don Simon (2004), "Bias at the Racetrack: Effects of Individual Expertise and Task Performance on Predecision Reevaluation of Alternatives," *Personality and Social Psychology Bulletin*, 7, 891-904.
- Brunnermeier, Markus and Jonathan Parker (2005), "Optimal Expectations," *American Economic Review*,
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Schleifer (2018), "Diagnostic Expectations and Credit Cycles," *Journal of Finance* 73, 199-227.
- Caplin, Andrew, and John Leahy (2001), "Psychological Expected Utility," *Quarterly Journal of Economics* 116, 55-79.
- Caplin, Andrew, and John Leahy (2004), "The Supply of Information by a Concerned Expert," *Economic Journal* 114, 487-505.
- Caplin, Andrew, and John Leahy (2006), "The Social Discount Rate," *Journal of Political Economy* 112, 1257-1268.
- Cooper, Arnold, Carolyn Woo, and William Dunkelburg (1988), "Entrepreneur's Perceived Chances of Success," *Journal of Business Venturing* 3, 97-108.
- Debondt, Werner, and Richard Thaler (1996), "Financial Decision Making in Markets and Firms: A Behavioral Perspective," *Handbook in Operations Research and Management Science* 9, North-Holland.

- Epstein, Larry, and Stanley Zin (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-969.
- Fuster, Andreas, David Laibson, and Brock Mendel (2012), "Natural Expectations and Economic Fluctuations," *Journal of Economic Perspectives* 24, 87-84.
- Gabaix, Xavier (2014), "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics* 129, 1661-1710.
- Hamilton, Barton (2000), "Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-employment," *Journal of Political Economy* 103, 604-631.
- Hansen, Lars, and Thomas Sargent (2008), *Robustness*, Princeton: Princeton University Press.
- Jevons, William (1905), *Essays in Economics*, London: Macmillan.
- Knox, R. E., Inkster, J. A. (1968), "Postdecision Dissonance at Posttime," *Journal of Personality and Social Psychology*, 18, 319-323.
- Lord, Charles, Lee Ross and Mark Lepper (1979), "Biased Assimilation and Attitudinal Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology*, 37, 2098-2109.
- Loewenstein, George (1987), "Anticipation and the Valuation of Delayed Consumption," *The Economic Journal*, 97, 666-684.
- Mijovic-Prelec, Canica, and Drazen Prelec (2010), "Self-Deception as Self-Signalling: A model and Experimental Evidence," *Philosophical Transactions of the Royal Society, B*, 265, 227-240.

6 Appendix

6.1 Derivation of (3)

Consider the maximization problem (2) for fixed $a \in A$. The first order condition for $p(\omega)$ is

$$0 = \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2} - \frac{1}{\theta} \ln \frac{p(\omega)}{p^*(\omega)} - \frac{1}{\theta} - \lambda$$

where λ is the Lagrange multiplier on the constraint $\sum_{\omega} p(\omega) = 1$. Note we do not have to consider the constraint $p(\omega) \geq 0$ as this constraint is never binding. Solving for $p(\omega)$

$$p(\omega) = p^*(\omega) \exp \left[\theta \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \theta \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2} - 1 - \theta\lambda \right] \quad (10)$$

Summing over ω implies

$$1 = \sum_{\omega} p(\omega) = \sum_{\omega} p^*(\omega) \exp \left[\theta \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \theta \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2} - 1 - \theta\lambda \right] \quad (11)$$

Dividing the left-hand side of (10) by the left-hand side of (11) yields,

$$p(\omega) = \frac{p^*(\omega) \exp \left[\theta \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \theta \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2} \right]}{\sum_{\omega} p^*(\omega) \exp \left[\theta \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \theta \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2} \right]}$$

Note that the $e^{-1-\theta\lambda}$ terms drop out.

To complete the derivation note,

$$\frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)} = \frac{\mu(\omega)u(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} - \sum_{\omega'} \frac{p(\omega')\mu(\omega')\mu(\omega)u(\omega')}{[\sum_{\omega''} p(\omega'')\mu(\omega'')]^2}$$

where E_{γ} is the expectation with respect to γ . Also note that given $\bar{p} \in \text{int}\Delta(\Omega)$, $p \in \text{int}\Delta(\Omega)$ so the constraint $p(\omega) \geq 0$ is never binding.

6.2 Derivation of the Value of action a

Substituting (3) into (2)

$$\begin{aligned}
& \sum_{\omega} \frac{p(\omega)\mu(\omega)}{\sum_{\omega'} p(\omega')\mu(\omega')} u(\omega) - \frac{1}{\theta} \sum_{\omega} p(\omega) \ln \frac{p(\omega)}{\bar{p}(\omega)} \\
&= \sum_{\omega} \frac{\frac{\bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right]}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]}{\sum_{\omega'} \frac{\bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]}{\sum_{\omega''} \bar{p}(\omega'') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega'')}{\partial p(\omega'')}\right]} \mu(\omega) - \frac{1}{\theta} \sum_{\omega} \frac{\bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right]}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]} \ln \frac{\bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right]}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]}{\bar{p}(\omega)} \\
&= \frac{\sum_{\omega} \bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right] \mu(\omega) u(\omega)}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right] \mu(\omega')} - \frac{1}{\theta} \frac{\sum_{\omega} \bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right]}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]} \ln \frac{\exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right]}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]} \\
&= \frac{\sum_{\omega} \bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right] \mu(\omega) u(\omega)}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right] \mu(\omega')} - \frac{\sum_{\omega} \bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right] \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]} + \frac{1}{\theta} \ln \sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]
\end{aligned}$$

look at the term

$$\begin{aligned}
& \frac{\sum_{\omega} \bar{p}(\omega) \exp\left[\theta \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}\right] \frac{\partial E_{\gamma} u(\omega)}{\partial p(\omega)}}{\sum_{\omega'} \bar{p}(\omega') \exp\left[\theta \frac{\partial E_{\gamma} u(\omega')}{\partial p(\omega')}\right]} \\
&= \sum_{\omega} p(\omega) \left(\frac{\mu(\omega)}{\sum_{\omega'} \mu(\omega') p(\omega')} (u(\omega) - E_{\gamma} u(\omega)) \right) \\
&= \sum_{\omega} \left(\frac{p(\omega)\mu(\omega)}{\sum_{\omega'} \mu(\omega') p(\omega')} (u(\omega) - E_{\gamma} u(\omega)) \right) \\
&= 0
\end{aligned}$$

6.3 Proof of the Propositions

Propositon 1 Given μ and \bar{p} , there exists a p that satisfies (1) for all $\omega \in \Omega$.

Proof: Define the mapping $T : \Delta(\Omega) \rightarrow \Delta(\Omega)$. Given any $p \in \Delta(\Omega)$, substitute into the right hand side of (1), and let $T(p)$ equal the left hand side. Clearly $T(p) \in \Delta(\Omega)$. Also, $T(p)$ is continuous in p . By Brouwer's fixed point theorem, there exists $p \in \Delta(\Omega)$ such that $T(p) = p$. ■

Proposition 2 With two states:

1. $p_H > \bar{p}_H$

2. p_H is increasing in $u_H - u_L$.
3. p_H is increasing in θ .
4. p_H is increasing in \bar{p}_H .
5. There exists $\bar{\mu}$ such that p_H is increasing in μ_H for $\mu_H < \bar{\mu}$ and decreasing in μ_H for $\mu_H > \bar{\mu}$.

Proof (Sketch): At any solution

$$p_H = \frac{\bar{p}_H}{\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}}$$

(1) follows from the observation that $u_H > u_L$ implies $e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}} < 1$.

Consider now T mapping defined in Proposition 1 and applied to the example with two states.

$$T(p_H) = \frac{\bar{p}_H}{\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_H \mu_L (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^2}}}$$

Note that when $p_H = 0$, we have

$$T(0) = \frac{\bar{p}_H}{\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_H}{\mu_L} (u_H - u_L)}} > 0$$

and when $p_H = 1$

$$T(1) = \frac{\bar{p}_H}{\bar{p}_H + \bar{p}_L e^{-\frac{\theta \mu_L}{\mu_H} (u_H - u_L)}} < 1$$

Now consider the derivative of the value function

$$\frac{p_H \mu_H u_H + p_L \mu_L u_L}{p_H \mu_H + p_L \mu_L} - \frac{1}{\theta} p_H \ln p_H / \bar{p}_H - \frac{1}{\theta} p_L \ln p_L / \bar{p}_L$$

which is

$$\frac{\mu_H u_H - \mu_L u_L}{p_H \mu_H + p_L \mu_L} - \frac{p_H \mu_H u_H + p_L \mu_L u_L}{(p_H \mu_H + p_L \mu_L)^2} (\mu_H - \mu_L) - \frac{1}{\theta} \ln p_H / \bar{p}_H + \frac{1}{\theta} \ln(1 - p_H) / \bar{p}_L$$

Note that this is positive and infinite if $p_H = 0$ or 1. Hence the first and last critical points are maxima. Generically every local maximum is associated with a point $T(p) = p$ such that $T'(p) < 1$.

(2), (3) and (4) follow from the observation that T is increasing in $\mu_H - \mu_L$, θ , and \bar{p}_H together with the observation that $T'(p) < 1$ at all maxima.

Consider the effect of an increase in μ_H on $\frac{\mu_H \mu_L}{(p_H \mu_H + p_L \mu_L)^2} = \frac{\mu_H(1-\mu_H)}{(p_H \mu_H + p_L(1-\mu_H))^2}$ the derivative is

$$\frac{1 - 2\mu_H}{(p_H \mu_H + p_L \mu_L)^2} - 2 \frac{(\mu_H - \mu_H^2)(p_H - p_L)}{(p_H \mu_H + p_L \mu_L)^3}$$

This has the same sign as

$$(1 - 2\mu_H)(p_H \mu_H + p_L \mu_L) - 2(\mu_H - \mu_H^2)(p_H - p_L) = 1 - \mu_H - p_H$$

So $\frac{\mu_H(1-\mu_H)}{(p_H \mu_H + p_L(1-\mu_H))^2}$ is increasing in μ_H when $\mu_H < 1 - p_H$ and decreasing otherwise. It follows that $T(p)$ is increasing in μ_H when $\mu_H < 1 - p_H$ and decreasing otherwise, which implies (4). ■

Proposition 3 With two states:

1. If $\mu_H > \mu_L$, then there is only one solution to (5).
2. Given μ and s , there exists $\bar{\theta}$ such that there is only one solution to (5).

Proof (sketch): There can be solution to (5) if the value function is strictly concave. the derivative of the value function is:

$$\frac{dV}{dp_H} = \frac{\mu_H \mu_L}{(p_H \mu_H + p_L \mu)^2} [u_H - u_L] - \frac{1}{\theta} \ln p_H / \bar{p}_H + \frac{1}{\theta} \ln p_L / \bar{p}_L$$

The second derivative is

$$\frac{d^2V}{dp_H^2} = -2 \frac{\mu_H \mu_L (\mu_H - \mu_L) (u_H - u_L)}{(p_H \mu_H + p_L \mu_L)^3} - \frac{1}{\theta} \frac{1}{p_H p_L}$$

This is negative if $\mu_H > \mu_L$, which establishes (1). The first term is continuous in p_H on the interval $[0, 1]$, and is therefore bounded. We can therefore choose, θ small enough that the expression is negative, which establishes (2).