NBER WORKING PAPER SERIES

# WISHFUL THINKING

Andrew Caplin
John V. Leahy

Working Paper 25707
http://www.nber.org/papers/w25707

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2019

We thank Anmol Bhandari, Stefano DellaVigna, Wouter den Haan, Behzad Diba, Jordi Gali, Nicola Pavoni, Kaitlin Raimi, Matthew Shapiro, Allen Sinai, Alp Simsek, Linda Tesar, Fabrizio Perri, Jaume Ventura, and Gianluca Violante for helpful discussions. We thank the Alfred P. Sloan and NOMIS Foundations for support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w25707.ack

Wishful Thinking
Andrew Caplin and John V. Leahy
NBER Working Paper No. 25707
March 2019
JEL No. D84,D91,E03,E71

**ABSTRACT**

We model agents who get utility from their beliefs and therefore interpret information optimistically. They may exhibit several biases observed in psychological studies such as optimism, procrastination, confirmation bias, polarization, and the endowment effect. In some formulations, they exhibit these biases even though they are subjectively Bayesian. We argue that wishful thinking can lead to reduced saving, can make possible information-based trade, and can generate asset bubbles.

Andrew Caplin
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

John V. Leahy
Gerald R. Ford School of Public Policy
University of Michigan
3308 Weill Hall
735 S. State St. #3308
Ann Arbor, MI 48109
and NBER
jvleahy@umich.edu

# Wishful Thinking*

Andrew Caplin[†] and John Leahy[‡]

March 2019

**Abstract**

We model agents who get utility from their beliefs and therefore interpret information optimistically. They may exhibit several biases observed in psychological studies such as optimism, procrastination, confirmation bias, polarization, and the endowment effect. In some formulations, they exhibit these biases even though they are subjectively Bayesian. We argue that wishful thinking can lead to reduced saving, can make possible information-based trade, and can generate asset bubbles.

"For a man always believes more readily that which he prefers" Francis Bacon (1620)

## 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

*We thank Anmol Bhandari, Stefano DellaVigna, Wouter den Haan, Behzad Diba, Jordi Gali, Nicola Pavoni, Kaitlin Raimi, Matthew Shapiro, Allen Sinai, Alp Simsek, Linda Tesar, Fabrizio Perri, Jaume Ventura, and Gianluca Violante for helpful discussions. We thank the Alfred P. Sloan and NOMIS Foundations for support.

†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

form expectations rationally. There is a lot of psychological evidence, however, that agents are poor information aggregators. They are often over-confident. They tend to interpret information in accordance with their priors. Beliefs seem to be sensitive to rewards. 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. Agents choose beliefs that raise their subjective utility subject to a cost. This desire to "see the world through rose colored glasses" naturally leads to several apparent deviations from rationality such as optimism, procrastination, confirmation bias, and polarization. We illustrate the economic implications in a series of simple examples and models. We show how wishful thinking may reduce saving and create bubbles. We show how wishful thinkers may engage in information based trade.

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 current wellbeing depends on the anticipation of future outcomes. Caplin and Leahy argue that 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." Jevons is even more direct. He argues that agents only care about the future because beliefs affect wellbeing in the present. In either case, the dependence of current wellbeing on beliefs about the future creates an incentive to believe that "good" outcomes are more likely than "bad" outcomes. As in the Bacon quotation above, 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 whatever one liked. Our model of the constraints rests on the idea that there are often lots of possible beliefs that are consistent with experience. Experience is rarely definitive. We limit our agents to "plausible beliefs", by which we mean beliefs that are not obviously contradicted by the available evidence. We do not model the technology by which agents choose beliefs. Instead 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 likelihood that the subjective beliefs would be rejected in favor of the objective ones. The more unlikely the truth, the more costly the beliefs.

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 also not magical thinking. Low probability events remain low probability, and zero probability events remain zero probability. Wishful thinking is strongest when outcomes are uncertain and payoff differences are large. Such situations might plausibly include choices that are made infrequently so that the agent lacks experience. Planning for retirement is an example. 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]

We present two formulations of the cost function which differ as to how beliefs are chosen. In the first formulation, we allow agents to directly choose their posterior beliefs. This formulation is simpler, and allows us to directly incorporate many of the tools developed by Hansen and Sargent (2008) to study robust control. It leads directly to optimism and over-confidence as agents overstate the probability of desired states. Wishful thinkers, in this formulation, tend to value uncertainty as it opens up the possibility of wishful thinking. This can lead them to delay actions when the future is uncertain, which can explain procrastination. It also can lead them to take actions that are themselves uncertain, which can help explain bubbles.

In the second formulation, we allow agents to choose how they interpret signals which then enter into their posterior beliefs. This formulation is mathematically more complex, but gives rise to a richer set of apparent deviations from rationality. In addition to optimism and procrastination, agents appear to twist information in the direction of their priors, a phenomenon known in psychology as confirmation bias. After observing the same information, two agents with divergent priors may also come to hold these divergent beliefs more strongly, a phenomenon known as polarization. This second formulation has the additional

[1] Reinhart 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.

distinction that these apparent deviations from rationality all occur in a model in which agents are subjective Bayesians. In this formulation, agents accurately combine subjective signals with priors to form subjective posteriors, so that their subjective reality is rational. Their problem is that is that their subjective interpretation of the signals differs from the objective counterpart. While subjectively Bayesian, they will appear non-Bayesian to an objective observer. We label the first formulation the cumulative model since the cost is on the posterior, which is the sum of all accumulated information, and we label the second formulation the flow model since the cost is on the flow of new information.

We present three economic implications of wishful thinking. First, we extend the cumulative model to a dynamic setting and present a simple model of consumption and saving in the presence of idiosyncratic income risk in the style of Huggett (1993). The economy is populated by both wishful thinkers and objective agents. Wishful thinkers tend to be optimistic regarding their future labor income. They therefore tend to place relatively high weight on high utility states which also tend to be the low marginal utility states. This leads them to consume more and save less, and therefore accumulate less wealth than do the objective agents. Second, we show that our setting is outside of the class of models considered by Milgrom and Stokey (1982). We present an example in which a wishful thinker and an objective agent actively trade based on private information. In the example, agents do not hold their beliefs dogmatically. They learn from the other agent's desire to trade. Nevertheless they agree to disagree. The wishful thinker knows the beliefs of the objective agent but chooses to believe differently. Third, we sketch a model of asset bubbles. We consider the introduction of a new and uncertain technology and argue that the wishful thinkers will bid its price above its fundamental value, where fundamental value is determined by the valuation of the objective agents. We map our description of a bubble to the narrative in Kindelberger (1978).

Our premise is that people shade their beliefs in ways that make desired states more likely. 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 higher odds on their preferred horse when interviewed after placing their bets than bettors did when interviewed while in line 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 controlled experimental 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. Closer to economics, Ito (1990) surveyed the exchange rate expectations of Japanese firms and found evidence of "wishful expectations." Firms expected the exchange rate to change in directions that benefitted them. Exporters tended to expect greater yen depreciation, whereas importers expected greater yen appreciation. More recently, Exley and Kessler (2018) find that agents interpret uninformative signals of their ability as positive signals, a phenomenon that they attribute to motivated reasoning.

While we consider the effect of anticipatory utility on belief choice, there may be other motivations for choosing beliefs that differ from objective reality. Beliefs may aid in goal attainment (Benabou and Tirole, 2002). People may react to the fear disappointment (Loomes and Sugden, 1982) or a concern for robustness (Hansen and Sargent, 2008). People might fear that their beliefs will affect outcomes (Pronin et al., 2006). Some of these alternative theories may predict people choose pessimistic beliefs. It is not our contention that all people at all times choose beliefs that make them happier. Rather we believe that wishful thinking is an important component of belief choice that can help to explain a large variety of apparently irrational behavior.

Section 2 discusses related literature. Section 3 presents the cumulative model. Section 4 presents the flow model. Section 5 discusses the economic applications. Section 6 discusses a number of issues, including alternative modeling choices and the relationship to the literature on robust control. Section 7 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). Benabou and Tirole 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. In this latter class, 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 contribution to this literature. 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 because it reduces fear, and fear reduces utility in their model. 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 they believe that the risk of an accident is low. Akerlof and Dickens assume that the agent chooses beliefs balancing the gain in belief utility against the cost of suboptimal decisions. The agent uses the objective probabilities when making this choice. Brunnermeier and Parker (2005) is another closely related paper. They assume that an agent chooses their beliefs at the beginning of their life prior to making any other decisions. Given this prior, the agent then behaves as a Bayesian in all subsequent periods. Like Akerlof and Dickens, they model belief choice as balancing the gain to anticipating a more positive future against the cost of suboptimal decisions. Like Akerlof and Dickens the agent evaluates belief choice using the objective probability distribution, and then proceeds with the chosen beliefs. There is a sense in which both of these models are models with divided selves.[2] When choosing beliefs the agent evaluates outcomes using the objective probabilities. When choosing actions the agent uses the chosen beliefs to evaluate the same outcomes. This tension makes it difficult to place these models into dynamic settings in which beliefs are chosen repeatedly over time, since it is difficult to model an agent who switches from being objective in one setting to optimistic in another. In our model, the only role of the objective beliefs is to anchor the cost of distorting beliefs. The objective beliefs are not used to evaluate the costs and benifits of choices. We discuss the relationship between our model and Brunnermeier and Parker's in detail in Section 6 below.

A second literature is the literature on anticipatory utility. Jevons (1905), Loewenstein (1987) and Caplin and Leahy (2001, 2004) all suppose that current happiness depends in some way on beliefs regarding future outcomes. Jevons believed that agents acted only to maximize current happiness. Intertemporal optimization, in this view, maximized the sum of the happiness from actions today and 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.

Our paper also contributes to the recent explosion of work that deviates from rational expectations. A partial and incomplete selection includes the following. Hansen and Sargent (2008) consider robust expectations that incorporate a fear of model misspecification. 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

[2] One interpretation of the Brunnermeier and Parker model is that parents choose beliefs for their children. In this interpretation there really are two selves.

weight on variables that are of greater importance. Burnside, Eichenbaum and Rebelo (2016) develop a model in which the distribution of expectations in the economy are influenced by social dynamics. 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. All of these papers focus on the consumers of information. Mullainathan and Schleifer (2005) model the supply side. They show how the public's preference for biased information affects the supply of information that is produced by the news media.

# 3 A Simple Model of Belief Choice

We begin with the cumulative information model in which the agent chooses their posterior beliefs and discuss the flow model in the next section. The essential elements of both theories 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 $\omega$ from a finite set of potential states of the world $\Omega$. There is an objective probability distribution over the second-period states $\bar{\gamma} \in \text{int}\Delta(\Omega)$, which we associate with the beliefs of an objective observer.

We follow Jevons (1905) and assume that the agent maximizes their current subjective expected utility. Current subjective expected utility incorporates both utility from current experience and utility from the anticipation of future outcomes. For simplicity we abstract from the former and focus solely on the anticipation of the future. We break down the utility from the anticipation into two components. The first is the payoff that the agent anticipates receiving in state $\omega$ should they choose action $a$, which we denote $u(a, \omega)$. The second is the agents' subjective probability that state $\omega$ will occur, which we denote $\gamma(\omega)$. We assume the agent understands their preferences so that they have an accurate assessment of $u(a, \omega)$.[3] We allow the agent's subjective beliefs $\gamma$ to differ from $\bar{\gamma}$. Subjective beliefs $\gamma$ may also depend

[3]Nothing is lost in assuming that the agent cannot manipulate their beliefs regarding $u(a, \omega)$. We can interpret $\omega$ broadly as including the quality of the match between the agent and the action $a$. An increase in the probability of a good match is equivalent to an increase in $u(a, \omega)$.

on the action choice $a$. The agent's subjective expected utility from the action $a$ is:[4,5]

$$\sum_{\omega \in \Omega} \gamma(\omega) u(a, \omega) \tag{1}$$

Maximizing (1) generates an incentive to choose beliefs. We discipline this choice with a cost of distorting beliefs. 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.[6] We follow Hansen and Sargent (2008) and relate the cost of choosing $\gamma(\omega)$ to the Kullback-Leibler divergence from $\bar{\gamma}(\omega)$ to $\gamma(\omega)$. The cost is:

$$\frac{1}{\theta} \sum_{\omega \in \Omega} \gamma(\omega) \ln \frac{\gamma(\omega)}{\bar{\gamma}(\omega)} \tag{2}$$

This cost (2) is the expected likelihood ratio under the subjective measure $\gamma$. It measures the ability of the agent to discriminate between $\gamma$ and $\bar{\gamma}$ given that the agent believes $\gamma$. The idea is that it is easier to choose a subjective belief that is not wildly contradicted by experience.

This cost function provides the link between the agent's subjective reality $\gamma$ and the outside world $\bar{\gamma}$. We will think of $\bar{\gamma}$ as reflecting the objective view of an unbiased expert. Agents understand what the expert is saying, but are aware that experts are not all knowing and are frequently wrong. Even in the best of cases, estimates of models come with standard errors. The cost function states that beliefs become harder and harder to justify, the further they are from expert opinion.

An alternative interpretation of the model is there exists a collection of experts with differing views. Such would be the case with many fundamental macroeconomic questions such as size of the fiscal multiplier, the slope of the Phillips curve, or the level of the natural rate of unemployment. In these cases there is no consensus on the true structural model and

[4] The advantage of Jevon's formulation is that the objective probabilities do not enter this calculation. The agent does not directly care about their future self and therefore does not consider the potential costs of mistaken beliefs as they do in the models of Akerlof and Dickens (1982) and Brunnermeier and Parker (2005). An interesting way forward is incorporate a concern for future mistakes into the model.

[5] The agent receives utility from the action $a$ both through prior anticipation and eventual experience. In Jevons' view, 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 invloves additional complications. See Caplin and Leahy (2006).

[6] See Benebou and Tirole (2016) for a discussion of theories of selective memory, selective attention and self-signalling.

experts come to vastly different conclusions. In this interpretation, $\bar{\gamma}$ represents mainstream opinion. The agent chooses which expert to believe and sees an increasing cost to choosing an expert that deviates too far from the consensus. In this interpretation, $\bar{\gamma}$ need not reflect the "truth" or "objective opinion." Instead $\bar{\gamma}$ might reflect the prevailing orthodoxy and the cost function might reflect the cost of deviating from this orthodoxy.[7]

The parameter $\theta$ captures the ease with which the agent can manipulate their beliefs. The larger is $\theta$ the greater the amount of evidence the agent would need before they reject their chosen beliefs in favor of the objective ones. When $\theta$ is equal to infinity, any beliefs are possible. When $\theta$ is equal to zero, any deviation from $\bar{\gamma}$ comes at an infinite cost.

Summarizing the above, the agent's maximization problem becomes:

$$V(\bar{\gamma}) = \max_{\gamma \in \operatorname{int}\Delta(\Omega), a \in A} \sum_{\omega \in \Omega} \gamma(\omega) u(a, \omega) - \frac{1}{\theta} \sum_{\omega \in \Omega} \gamma(\omega) \ln \frac{\gamma(\omega)}{\bar{\gamma}(\omega)}. \tag{3}$$

The first term is the subjective expected utility of choice $a$ given the belief $\gamma$. The second term is the cost of the believing $\gamma$. Implicit in the maximization problem (3) is the assumption that the agent understands that their beliefs $\gamma$ will depend on the choice $a$. We consider the implications of assuming that the agent is naïve in Section 6.

The maximization problem (3) is very similar to the robust control problem in Hansen and Sargent (2008). There are two differences. First, Hansen and Sargent maximize over $a$ and conditional on $a$ minimize over $\gamma$. Second, because Hansen and Sargent minimize over $\gamma$, they add rather than subtract the cost (2). This is equivalent to replacing $\theta$ with $-\theta$. It is therefore not surprising that many of our conclusions will be exactly the opposite of those of Hansen and Sargent.

## 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 $\gamma(\omega)$ implies

$$\gamma(\omega) = \frac{\bar{\gamma}(\omega) \exp[\theta u(\omega)]}{\sum_{\omega' \in \Omega} \bar{\gamma}(\omega') \exp[\theta u(\omega')]} \tag{4}$$

[7] If we take this view then an interesting way forward is to endogenize the supply of experts along the lines of Mullainathan and Schleifer (2005) .

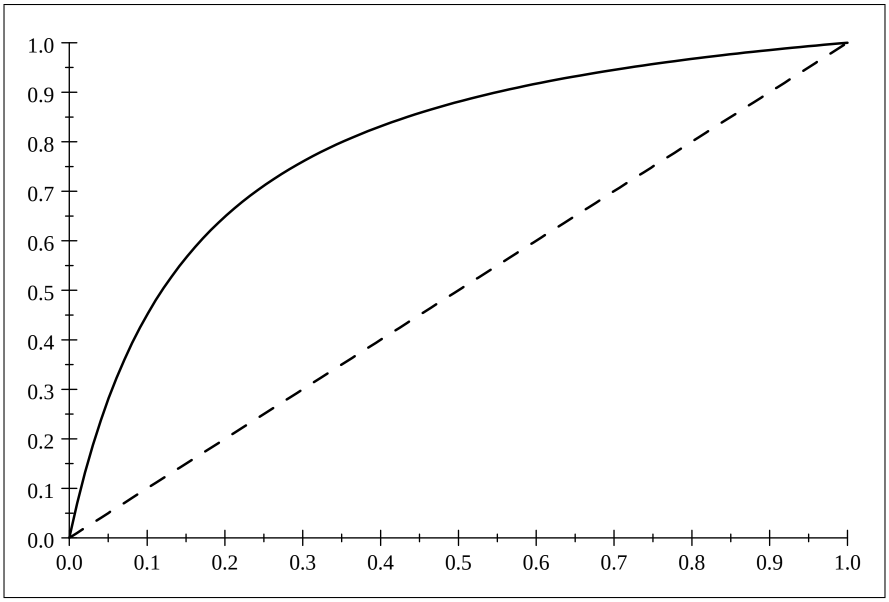


Figure 1: The figure depicits $\gamma_H$ as a function of $\bar{\gamma}_H$ for $\mu_L = \mu_H = .5$ and $\theta(u_H - u_L) = 2$.

According to (4), if $u(\omega) = u$ for all $\omega$, then the agent has no incentive to distort beliefs.[8] Otherwise, the agent distorts their beliefs. They tend to increase the probability of states with high utility and they reduce the probability of states with low utility.

Figure 1 graphs $\gamma(\omega)$ as a function of $\bar{\gamma}(\omega)$ for an example with two states: a high utility state $\omega_H$ in which utility is $u_H$ and a low utility state $\omega_L$ with utility $u_L < u_H$. The objective probability of the high utility state is on the horizontal axis. The solid line represents the chosen probability as measured on the vertical axis. The dashed line is the 45 degree line. The gap between the two lines represents the extent of wishful thinking,

$$\gamma_H - \bar{\gamma}_H = \frac{\bar{\gamma}_H \bar{\gamma}_L \left(\exp\left[\theta u_H\right] - \exp\left[\theta u_L\right]\right)}{\bar{\gamma}_H \exp\left[\theta u_H\right] + \bar{\gamma}_L \exp\left[\theta u_L\right]}$$

Since the $\omega_H$ is the good state, $\exp\left[\theta u_H\right] - \exp\left[\theta u_L\right] > 0$ and 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 is not magical thinking. If $\bar{\gamma}_H$ is zero, then $\gamma_H$ is zero too. Once the US team has been eliminated from the world cup, it is difficult to for Americans to fantasize about winning the tournament.

[8] Exley and Kessler (2018) find that when they remove motivation for distorting beliefs, agents appear unbiased.

Wishful thinking is strongest when the state is unknown. 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.[9]

It is easy to show that the difference between $\gamma_H$ and $\bar{\gamma}_H$ peaks at $\bar{\gamma}_H < \frac{1}{2}$, so that wishful thinking is strongest when the good state $\omega_H$ is slightly less likely than the bad state. This has two implications. First, the agent interprets uninformative signals positively. Consistent with this result, Exley and Kessler (2018) find that subjects in their experiment update favorably upon receiving a signal that is known to be uninformative. Second, there is a range of $\bar{\gamma}_H$ for which $\bar{\gamma}_H < \frac{1}{2}$ and $\gamma_H > \frac{1}{2}$, so that the wishful thinker believes that the good state is more likely than the bad state, whereas the objective observer believes the opposite. 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 $\bar{\gamma}_H$ must stay close enough to fair, so that agents can believe $\gamma_H > \frac{1}{2}$.

## 3.2 Action Choice

We can calculate the value of the action $a$ under the optimal beliefs. Substituting (4) into $V(\bar{\gamma})$ for a given action choice $a$ yields

$$V(\bar{\gamma}) = \max_a \frac{1}{\theta} \left( \ln \sum_{\omega} \bar{\gamma}(\omega) e^{\theta u(a,\omega)} \right) \tag{5}$$

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.

While we have modeled an agent who distorts their beliefs, choice in this model turns out to be observationally equivalent to the choice of an agent with distorted utility. One way to test for the difference between these two settings would be to combine data on choice and beliefs.

[9] One has to be careful in defining the state space $\Omega$. For example, the outcome of the flip of a fair coin is unknown, but there may be little room for wishful thinking if it is known that the coin is fair. In such a case $\theta$ is very small and it is difficult to believe $\gamma$ is too different from $\bar{\gamma}$ Alternatively, one might take the state to be the fairness of the coin, in which case wishful thinking is more likely to be significant.

## 3.3 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 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 $\omega_1$ through $\omega_N$ equally spaced around a circle, so that the distance between $\omega_1$ and $\omega_2$ is equal to the distance between $\omega_1$ and $\omega_N$. Nature picks the true state $\hat{\omega} \in \Omega$, and the agent picks a state $a \in \Omega$. 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 $\delta$: $u(a, \hat{\omega}) = u(\delta(a, \hat{\omega}))$. This payoff function captures both types of experiment. In the first case, there is a correct answer and a collection of incorrect answers: $u(\delta) = 1$ if $\delta = 0$ and $u(\delta) = 0$ otherwise. In the second case, the loss is increasing in the $\delta$. Suppose that the objective expert has a uniform prior over the states and receives a signal $s \in \Omega$ that has an objective density that is symmetric about the true state and declining in the distance from the true state. The symmetry of the prior and the signal imply that the experts beliefs are symmetric about the signal: $\bar{\gamma}(\omega|\omega = s + x) = \bar{\gamma}(\omega|\omega = s - x)$ for all $x \leq |\Omega|/2$.

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. The first-order condition (4) becomes:

$$\gamma(\omega) = \frac{\bar{\gamma}(\omega) \exp\left[\theta u(\delta(s, \omega))\right]}{\sum_{\omega'} \bar{\gamma}(\omega') \exp\left[\theta u(\delta(s, \omega))\right]}. \tag{6}$$

Without loss of generality label the signal state $\omega_0$. Label the states to the right of $\omega_0$ (as we move about the circle): $\omega_1$, $\omega_2$,..., and label the states to the left $\omega_{-1}$, $\omega_{-2}$,... 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 $\omega_0$, $\omega_{N/2}$. With this labeling, the absolute value of the index is equal to $\delta(s,\omega)$. Now consider the ratio implied by (6):

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

Several observations follow immediately. First, $\bar{\gamma}(\omega_n) = \bar{\gamma}(\omega_{-n})$ implies $\gamma(\omega_n) = \gamma(\omega_{-n})$, so that $\gamma$ inherits the symmetry of $\bar{\gamma}$. Second, since $u$ is maximized at $\delta = 0$, $\frac{\gamma(\omega_0)}{\gamma(\omega_n)} > \frac{\bar{\gamma}(\omega_0)}{\bar{\gamma}(\omega_n)}$ for all $\omega_n \neq \omega_0$. As the $\gamma$ sum to one, it follows that $\gamma(\omega_0) > \bar{\gamma}(\omega_0)$. The agent is overconfident that they have selected the correct state. Finally, if we consider any subset of states relatively close to $\omega_0$, $\bar{\Omega} = \{\omega_n \text{ s.t. } |n| < \bar{N}\}$, $\frac{\gamma(\omega_n)}{\gamma(\omega_m)} > \frac{\bar{\gamma}(\omega_n)}{\bar{\gamma}(\omega_m)}$ 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 overconfident in the sense that their subjective confidence intervals will be too tight.

## 3.4 Procrastination and Self-Control

An example illustrates how the model generates procrastination. Consider the decision problem in Figure 2. There are three periods represented by the nodes **A**, **B**, and **C**. In the first period (node **A**), the agent chooses whether to take an action or delay. The cost of the action is 1. If the agent delays, then the agent is either free or busy in period 2 (node **B**). If the agent is free, they take the action in period 2 at a cost of 1/2. If they are busy, they delay the action until period 3 (node **C**) and take the action at a cost of 2. The agent wishes to minimize costs.

[Insert Figure 2]

If we assume that it is equally likely that the agent is free or busy in period 2, then an objective agent would calculate the expected cost of delay in period 1 to be 5/4. The objective agent would then choose to act in period 1. What would a wishful thinker do? Conditional on delay in period 1, the wishful thinker would want to increase the probability of being free in period 2. If we take $\theta = 1$, then a direct application of (4) implies that the wishful thinker will anticipate an 82% chance of being free. Their subjective cost of delay in period 1 would then be .77. The cost of taking the action is still 1. The wishful thinker would then choose to delay. The agent overestimates their ability to complete the task in period 2 and puts off taking the action.

The reason that the agent procrastinates is that delay is the uncertain option and wishful thinkers value uncertainty. In a model of firm entry such as Dixit (1989), in which the value of delay is uncertain, wishful thinking will tend to reduce entry. In a model of firm entry such as Hopenhayn (1982), in which the value of entry is uncertain, wishful thinking will tend to increase entry.

A slight reformulation of the example fits the experimental evidence of DellaVigna and Malmendier (2006). They find that agents who purchase monthly gym memberships would save money if they instead paid for each visit separately. Moreover, they find that agents who purchase monthly memberships, which automatically roll-over in their data set, tend to cancel less often than agents who purchase annual memberships, which do not roll over. Their preferred explanation is that agents both overestimate their ability to attend the gym and overestimate their ability to cancel their membership when desired. In our example, delay would be analogous to the purchase of a monthly membership, and the second period action would be analogous to attendance or cancellation. Note that in our explanation, agents are time consistent. They correctly anticipate what they will do in each state of the world. Their mistake is that they endogenously overestimate the probabilities of the states in which they go to the gym and cancel their membership.

Another reformulation illustrates the planning fallacy of Kahneman and Tversky (1979). The planning fallacy is the tendency of people to underestimate the time that it will take to complete a task. To capture the planning fallacy, we relabel "free" as "short completion time" and "busy" as "long completion time" in the example. Agents will then choose to underestimate the time of completion because they want to believe that the completion time is small. There are many explanations of the planning fallacy in the psychology literature, but Buehler, Griffin and McDonald (1997) attribute it to wishful thinking.[10] In their experiments, they manipulate the incentive for early completion and show that incentives to complete a task earlier exacerbate the planning fallacy. Participants' beliefs respond to incentives.

Agents underestimate the time it will take to complete a task even though they have past experience with the task (Buehler, Griffen and Ross, 1994). This is also consistent with our theory. In our theory past experience would play a role similar to objective information $\bar{\gamma}$. So long as past experience is not definitive, there is room for wishful thinking. Buehler, Griffen and Ross discuss several reasons that agents might convince themselves that past experience is not definitive, including arguments such as *past experience is not comparable*, that *the situation has changed*, and that there might have been *extenuating circumstances.*

[10] Brunnermeier, Papakonstantinou, and Parker (2005) use the model of Brunnermeier and Parker (2005) to explain the planning fallacy.

# 4 The flow model[11]

In Section 3, we placed the information cost on the agent's posterior. The agent could manipulate their stock of information. An alternative approach is to place the cost on the flow of new information. A feature of this approach is that the agent may be subjectively Bayesian, yet exhibit non-Bayesian behavior to an outside observer.[12] They accurately combine their subjective interpretation of signals with their subjective priors to form subjective posteriors. Their subjective reality is therefore rational. Their bias is that they view the world optimistically and twist the interpretation of the signals that they observe. They see the world through rose colored glasses.

We alter the model of Section 3 to place the cost of information on the flow of new information. As before there is a set of actions $a \in A$, a set of future states $\omega \in \Omega$, and a payoff to each action in each state $u(a, \omega)$. We drop the objective posterior $\bar{\gamma}(\omega)$. Instead we assume that the agent has a prior over $\Omega$ and observes a signal which is informative about the realization of $\omega \in \Omega$. We denote agent's prior $\mu(\omega) \in \text{int}\Delta\Omega$. This prior need not equal the prior held by an objective observer, as it may be distorted by the interpretation of information in the past. The signal $s$ is drawn from some finite space of signals $S$. $s$ is generated by an information structure $\bar{p} : \Omega \to \Delta(S)$. An objective observer understands this information structure. The wishful thinker can choose to believe that the signal is generated by some alternative information structure $p : \Omega \to \Delta(S)$. The cost of choosing $p$ different from $\bar{p}$ is:

$$\frac{1}{\theta} \sum_{\omega \in \Omega} \sum_{s \in S} p(s|\omega) \ln \frac{p(s|\omega)}{\bar{p}(s|\omega)} \tag{7}$$

The agent distorts the mapping between states and signals. The idea is that after seeing a signal, the agent has an incentive to exaggerate the likelihood that it was generated by a desirable state and downplay the likelihood that it was generated by an undesirable state. We assume that the cost is independent of the prior $\mu$. Alternatively, we could have assumed that each term in (7) is weighted by $\mu(\omega)$ with the idea that it is more costly to twist probabilities in more likely states. The cost function (7) is simpler and leads to qualitatively similar phenomenon.

[11] We thank Nicola Pavoni and Alp Simsek whose comments greatly improved this section.

[12] The agents in Section 3 are not subjective Bayesians. New information alters the set of permissible beliefs leading agents to reinterpret past information as well.

The value to choosing an action $a \in A$ given the prior $\mu$ and the signal $s$ is then:

$$V(a,\mu,s) = \max_p \sum_{\omega\in\Omega} \frac{p(s|\omega)\mu(\omega)}{\sum_{\omega'} p(s|\omega')\mu(\omega')} u(\omega,a) - \frac{1}{\theta}\sum_{\omega\in\Omega}\sum_{s'\in S} p(s'|\omega)\ln\frac{p(s'|\omega)}{\bar{p}(s'|\omega)} \tag{8}$$

The first term is subjective expected utility. The agent combines their subjective interpretation of the signal with their prior according to Bayes rule. As the signal $s$ is known, expected utility depends only on the probability of observing the signal $s$ in each state. Note that there is no presumption that $\sum_{\omega\in\Omega} p(s|\omega) = 1$, as each $p(s|\omega)$ represents the probability of observing $s$ in a different state. The second term is the cost of information. This depends not only on the signal $s$, but the probability of observing other signals $s' \neq s$.

The agent chooses $a \in A$ to maximize (8).

## 4.1 Implications for belief choice

As before we fix the action $a$ and write $u(\omega)$ for $u(a,\omega)$. We also fix $s \in S$, and write $p(\omega)$ and $\bar{p}(\omega)$ for $p(s|\omega)$ and $\bar{p}(s|\omega)$ respectively. Finally let $\gamma(\omega)$ denote the posterior resulting from the choice of $p(\omega)$ :

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

The first order condition for $p(\omega)$ implies:

$$p(\omega) = \frac{\bar{p}(\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]}{\bar{p}(\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right] + (1-\bar{p}(\omega))} \tag{9}$$

where $E_\gamma u(\omega)$ is the expectation of $u(\omega)$ with respect to the subjective posterior $\gamma$, and $\frac{\partial E_\gamma u}{\partial p(\omega)}$ is the partial derivative of this expectation with respect to $p(\omega)$. The derivation of (9) is in the appendix.

According to (9), the agent distorts their interpretation of the signal whenever $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}$ is not zero. In this case, they tend to increase the probability of that the signal they received comes from a state if increasing that probability increases expected utility. The derivative can be written as

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

According to the term in brackets, the agent tends to raise the probability of states with above average utility (according to the posterior $\gamma$). 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 (9) is multiplied by $\theta$. The larger is $\theta$, the easier it is for the agent to manipulate their interpretation of the signal.[13]

Note that since $p$ affects $\gamma$, it enters both sides of (9). These first-order conditions therefore define $p$ implicitly. The next proposition shows that (9) is not vacuous. A solution always exists. All proofs are in the appendix.

**Proposition 1** Given $\mu \in \text{int}\Delta(\Omega)$ and $\bar{p} \in \text{int}[0,1]^{|\Omega|}$, there exists a $p \in \text{int}[0,1]^{|\Omega|}$ that satisfies (9) for all $\omega \in \Omega$.

It is possible that there are multiple solutions to (9). This can happen when $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}$ is increasing in $p(\omega)$, so that 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 $V(a, \mu, s)$.

Given that the flow model leads to equations that are less familiar than those of the cumulative model, it is useful to consider an example with two states to see the how the model works.

## 4.2 Example with two states

Suppose that there are two states $\omega_H$ and $\omega_L$ with $u_H \equiv u(\omega_H) > u(\omega_L) \equiv u_L$ so that $\omega_H$ is the good state and $\omega_L$ is the bad state. Let $\mu_H$ denote the prior belief that the state is $\omega_H$ and $\mu_L$ the prior that the state is $\omega_L$. Let $\bar{p}_H = \bar{p}(s|\omega_H)$ denote the objective probability that the signal is $s$ given that the state is $\omega_H$. Define $\bar{p}_L$, $p_H$, and $p_L$ accordingly.

With these definitions (9) can be written as,

$$\begin{aligned} p_H &= \frac{\bar{p}_H}{\bar{p}_H + (1-\bar{p}_H)e^{-\frac{\theta\mu_H\mu_L p_L(u_H-u_L)}{(p_H\mu_H+p_L\mu_L)^2}}} \\ p_L &= \frac{\bar{p}_L}{\bar{p}_L + (1-\bar{p}_L)e^{\frac{\theta\mu_H\mu_L p_H(u_H-u_L)}{(p_H\mu_H+p_L\mu_L)^2}}}. \end{aligned} \tag{10}$$

[13] If the terms in (7) are weighted by $\mu(\omega)$, then it is more costly to distort beliefs when the probability of the state is high. This effect offsets the desire coming from preferences to distort these beliefs.

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

**Proposition 2** With two states:

1. $p_H > \bar{p}_H$ and $p_L < \bar{p}_L$
2. $p_H$ is strictly increasing in $u_H - u_L$, whereas $p_L$ is strictly decreasing in $u_H - u_L$
3. $p_H$ is strictly increasing in $\theta$, whereas $p_L$ is strictly decreasing in $\theta$
4. $p_H$ is strictly increasing in $\bar{p}_H$, whereas $p_L$ is strictly increasing in $\bar{p}_L$
5. $p_H$ is strictly increasing in $\mu_H$ if $p_H\mu_H < p_L\mu_L$ and decreasing in $\mu_H$ if $p_H\mu_H > p_L\mu_L$. The opposite applies for $p_L$.

Point (1) is the essence of wishful thinking: the subjective probability that the signal originated from the good state rises, whereas the subjective probability that the signal originated from the bad state falls. The subjective interpretation of the signal is therefore more optimistic than the objective interpretation, which implies that the subjective posterior will also be more optimistic. Point (1) follows immediately upon inspection of the sign of the exponent in the exponential term in the denominator of (10). 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. The only surprising result is (5). To understand this result, note that $\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} = 0$ both when $\mu_H = 0$ and when $\mu_H = 1$, so $p_H = \bar{p}_H$ in each of the extreme cases. In between, $p_H$ rises and then falls relative to $\bar{p}_H$ as $\mu_H$ rises. It turns out that the inflection point is equal to $p_L\mu_L/\mu_H$.

### 4.2.1 Action Choice

We can calculate the value of the action $a$ under the optimal beliefs (see the appendix). Substituting (9) into $V(a, \mu, s)$ for a given action choice $a$ yields,

$$V(\mu, s) = \max_{a \in A} E_{\gamma(a)} u(a, \omega) + \frac{1}{\theta} \sum_{\omega \in \Omega} \ln \left( \bar{p}(\omega) \exp \left[ \theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} \right] + (1 - \bar{p}(\omega)) \right)$$

where $\gamma(a)$ is the posterior associated with action $a$.

Like the cumulative model, the flow model can produce overconfidence and procrastination. The flow model also leads to other apparent deviations from rational behavior such as confirmation bias and polarization.

## 4.3 Confirmation bias

Confirmation bias occurs when an agent interprets information in a way that conforms to 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 also likely correlated with payoffs if payoffs are persistent.

To illustrate confirmation bias, consider the example with two states, $\omega_1$ and $\omega_2$. Consider two agents, one of whom receives high utility from state $\omega_1$ and the other receives high utility in state $\omega_2$. 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 of the signal 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) considered a population of parents with similar priors: all professed to believe that home care is superior to day care for their children. They differed, however, in their payoffs, as some had chosen home care for their children, while others had chosen day-care. They found that the interpretation of evidence aligned with the payoffs of the subjects rather than the prior. The parents who had placed their children in day care rated a study supporting day care more favorably, whereas the parents who cared for their children at home did the opposite. In both of these studies, the interpretation of information appears to have been 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.

## 4.4 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 (Lord, Ross and Lepper, 1979). 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 setting with two agents labeled $i$ and $j$ and two states labeled $\omega_1$ and $\omega_2$. Suppose that agent $i$ receives utility $u_H$ in state $\omega_1$ and $u_L$ in state $\omega_2$ and agent $j$ receives utility $u_H$ in state $\omega_2$ and $u_L$ in state $\omega_1$. 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 $\omega_1$, and agent $j$ places the same prior on state $\omega_2$. Each then sees the same signal $s$ generated by the information structure $\bar{p}$. Let $\bar{p}_1$ denote the probability of seeing $s$ in state $\omega_1$ and $\bar{p}_2$ the probability of seeing $s$ in state $\omega_2$.

Each interprets the signal according to (10). Since their payoffs and priors differ, their interpretation differs. Agent $i$ chooses

$$
\begin{aligned}
p_1(i) &= \frac{\bar{p}_1}{\bar{p}_1 + (1-\bar{p}_1)e^{-\frac{\theta\mu_H(i)\mu_L(i)p_1(i)(u_H-u_L)}{(p_1(i)\mu_1(i)+p_2(i)\mu_2(i))^2}}} \\
p_2(i) &= \frac{\bar{p}_2}{\bar{p}_2 + (1-\bar{p}_2)e^{\frac{\theta\mu_H(i)\mu_L(i)p_2(i)(u_H-u_L)}{(p_1(i)\mu_1(i)+p_2(i)\mu_2(i))^2}}}
\end{aligned}
$$

whereas agent $j$ chooses

$$
\begin{aligned}
p_1(j) &= \frac{\bar{p}_1}{\bar{p}_1 + (1-\bar{p}_1)e^{\frac{\theta\mu_H(j)\mu_L(j)p_1(j)(u_H-u_L)}{(p_1(j)\mu_1(j)+p_2(j)\mu_2(j))^2}}} \\
p_2(i) &= \frac{\bar{p}_2}{\bar{p}_2 + (1-\bar{p}_2)e^{-\frac{\theta\mu_H(j)\mu_L(j)p_2(j)(u_H-u_L)}{(p_1(j)\mu_1(j)+p_2(j)\mu_2(j))^2}}}
\end{aligned}
$$

The functional forms are similar. Only the mapping between payoffs, priors, and states is switched. In keeping with Proposition 2, agent $i$ twists their interpretation of the signal in the direction of their more preferred state, which is state $\omega_1$. They choose $p_1(i) > \bar{p}_1$ and $p_2(i) < \bar{p}_1$. Agent $j$ twists the signal in the opposite direction. They choose $p_1(j) < \bar{p}_1$ and $p_2(j) > \bar{p}_1$.

These choices give rise to the posteriors:

$$\gamma_1(i) = \frac{\mu_H}{\mu_H + (1 - \mu_H)\frac{p_2(i)}{p_1(i)}}$$

and

$$\gamma_1(j) = \frac{\mu_L}{\mu_H + (1 - \mu_H)\frac{p_2(j)}{p_1(j)}}$$

Polarization occurs if $\gamma_i(\omega_1) > \mu_H$ and $\gamma_j(\omega_1) < \mu_L$ which requires $\frac{p_2(i)}{p_1(i)} < 1 < \frac{p_2(j)}{p_1(j)}$. In this case both agents have observed the same signal and each has become more confident in their assessment of the state.

Now $\frac{p_2(i)}{p_1(i)} < 1$ if

$$\bar{p}_2\left(\bar{p}_1 + (1-\bar{p}_1)e^{-\frac{\theta\mu_H(i)\mu_L(i)p_1(i)(u_H-u_L)}{(p_1(i)\mu_1(i)+p_2(i)\mu_2(i))^2}}\right) < \bar{p}_1\left(\bar{p}_2 + (1-\bar{p}_2)e^{\frac{\theta\mu_H(i)\mu_L(i)p_2(i)(u_H-u_L)}{(p_1(i)\mu_1(i)+p_2(i)\mu_2(i))^2}}\right).$$

Similarly $\frac{p_2(j)}{p_1(j)} > 1$ if

$$\bar{p}_2\left(\bar{p}_1 + (1-\bar{p}_1)e^{\frac{\theta\mu_H(j)\mu_L(j)p_1(j)(u_H-u_L)}{(p_1(j)\mu_1(j)+p_2(j)\mu_2(j))^2}}\right) > \bar{p}_1\left(\bar{p}_2 + (1-\bar{p}_2)e^{-\frac{\theta\mu_H(j)\mu_L(j)p_2(j)(u_H-u_L)}{(p_1(j)\mu_1(j)+p_2(j)\mu_2(j))^2}}\right).$$

Since $u_H > u_L$, the signs of the exponents in the exponentials push the inequalities in the right direction. If $\bar{p}_1 = \bar{p}_2$, then both inequalities hold. Hence polarization is possible. 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 when the signal is relatively uninformative about the state and the desire to believe is large.

## 4.5 Comparing the two formulations

The cumulative model and the flow model each have their advantages and disadvantages. Both approaches can explain overconfidence and procrastination. In addition, the flow model can explain confirmation bias and polarization. The cumulative model cannot. The reason is that subjective posteriors are closely tied to objective posteriors in (4). News that raises $\bar{\gamma}(\omega)$ will tend to raise $\gamma(\omega)$ for all agents. The cumulative model leads to an algebraically simpler solution which may prove useful in dynamic applications.

There are other differences in the two approaches. Agents in the flow model are subjective Bayesians and appear non-Bayesian to an objective observer. On the other hand, agents in the cumulative model maximize (3). 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 flow model. They evolve with the flow of information. Beliefs can potentially change dramatically in the cumulative 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.

If the payoffs to states are stable, then beliefs in the flow model may diverge further and further from objective reality over time as the agent continues to twist new signals in the direction of the higher payoffs. This divergence may seem undesirable in some settings. It is possible that the true model is a combination of the two settings. The flow model is used on a day to day basis, but ever so often the agent takes stock of their world view and recalibrates their posteriors using the cumulative model.

# 5 Three asset pricing models

In this section we illustrate some of the equilibrium implications of the theory through three asset pricing models. The first is a model with a risk free bond and idiosyncratic income risk along the lines of Huggett (1993) extended to include agents that differ in their optimism. All else equal the more optimistic agents tend to consume more and save less, and therefore are less wealthy in steady state. The equilibrium interest rate is above that in an economy without optimism. The second is a model of trade with private information. Again we assume that agents differ in their optimism. We show that informed traders, whether optimistic or not, can profit from their private information by trading with agents with different beliefs. The third model is a model of bubbles in the spirit of Kindleberger (1978). Occasionally optimists get lucky; their wealth increases; and their influence on asset prices grows, raising asset prices for a time above what is warranted by objective observers.

## 5.1 A Huggett economy

Time is discrete. There is a single asset, a risk free bond, with a gross return $R$. There are a continuum of agents indexed by $i \in [0,1]$. Agents indexed by $i \in [0,\eta]$ are objective. They maximize

$$E\sum_{t=0}^{\infty}\beta^t u(c_t)$$

The remainder of agents are wishful thinkers. Wishful thinkers maximize the same utility function $u$ and have the same discount factor $\beta$, but choose their expectations. In both cases we assume that $u$ is increasing, concave, differentiable and displays decreasing absolute risk aversion, a class which includes constant relative risk aversion.[14]

Each period $t$, each agent $i$ receives an endowment $y_{it}$. $y_{it}$ takes one of $S$ values $\{y_1, \ldots y_S\} \equiv Y$. The realizations of $y_{it}$ are independent and identically distributed across individuals. The probability of $y_{it}$ is $p(y_{it})$.

The state of individual $i$ in period $t$ is $(A_i, y_i)$ where $A_i$ is their in period $t-1$ saving and $y_i$ is their current-period endowment income. We restrict $A_i > \phi$ where $\phi < 0$ and $y_1 \geq -(R-1)\phi$ so that the agent can always service their debts if saving is negative.

We need to extend (3) to a dynamic setting. We construct the dynamic analog of the cumulative model. We maintain the assumption that wishful thinkers are sophisticated and that they understand the relationship between action choice and beliefs. Given their initial state $(A_0, y_0)$, we assume that they choose a sequence of state contingent plans $A(y^t)$ subjective beliefs $\gamma(y^t)$ where $y^t$ is the history of endowment realizations through period $t$. Their maximization problem is:

$$\begin{aligned} V^w(A_0, y_0) \quad = \quad & \max_{\{\gamma(y^t), A(y^t)>\phi\}_{t>0, y^t \in Y^t}} \sum_{t=0}^{\infty}\beta^t \sum_{y^t\in Y^t}\gamma(y^t)u(RA(y^{t-1}(y^t)) + y_t - A(y^t)) \\ & -\frac{1-\beta}{\theta}\sum_{t=1}^{\infty}\beta^t\sum_{y^t\in Y^t}\gamma(y^t)\ln\frac{\gamma(y^t)}{\bar{\gamma}(y^t)} \end{aligned}$$

where $y^{t-1}(y^t)$ denotes the history through $y^{t-1}$ embedded in the history $y^t$. As before the first term is subjective expected utility. The second term is the cost of distorting beliefs. We allow the agent to manipulate beliefs at all horizons. We discount the distortion of future beliefs at the same rate as future utility. The $1-\beta$ in front compensates for the fact that distorting $\gamma(y^t)$ also tends to distort $\gamma(y^s)$ for all $s > t$, so that $\gamma(y^t)$ effectively enters the sequence problem with a weight $\beta^t/(1-\beta)$.

[14] We need asset accumulation to be bounded to use standard theorems. The only cases in the literature in which it is known that asset accumulation is bounded is the case of i.i.d. income shocks and decreasing absolute risk aversion (Schechtman and Escudero, 1977) or income that follows a two state Markov chain and constant relative risk aversion (Huggett, 1983).

If we assume subjective beliefs are consistent in the sense that the conditional expectations satisfy the laws of probability, $\gamma(y^t) = \gamma(y_s(y^t))\gamma(y^t|y_s(y^t))$ for all $0 < s < t$, then we can write the cost of distorting $y^t$ recursively:

$$\sum_{y^t \in Y^t} \gamma(y^t) \ln \frac{\gamma(y^t)}{\bar{\gamma}(y^t)} = \sum_{y_1 \in Y} \left[ \gamma(y_1) \ln \frac{\gamma(y_1)}{\bar{\gamma}(y_1)} + \gamma(y_1) \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y^t|y_1) \ln \frac{\gamma(y^t|y_1)}{\bar{\gamma}(y^t|y_1)} \right]$$

The cost of distorting $y^t$ is the cost of distorting $y_1$ plus the cost of distorting $y^t$ conditional on $y_1$. This allows us to write the wishful thinker's problem recursively (see the appendix for the details):

$$\begin{aligned} V^w(A, y) &= \max_{A' > \phi, \{\gamma(y')\}_{y' \in Y}} u(c) + \beta \sum_{y' \in Y} \gamma(y') V^w(A', y') - \frac{\beta}{\theta} \sum_{y' \in Y} \gamma(y') \ln \frac{\gamma(y')}{\bar{\gamma}(y')} \qquad (11) \\ &= \max_{A' > \phi} u(c) + \frac{1}{\theta} \ln E_{(A,y)} \exp\{\theta V^w(A', y')\} \end{aligned}$$

where it is understood that $c = RA + y - A'$ and the second equality comes from substituting the optimal beliefs as in (5). Note that choice is dynamically consistent in this setting. Even though, following Jevons, we think of the agent as choosing the entire future sequence of beliefs and actions to maximize their current subjective utility, the agent solves a problem of similar form in the future so that the subjective plans that the agents contemplates in one period become actual plans when future states are realized. The wishful thinker distorts the probabilities that future states will occur but not the actions that will be taken in those states.

An equilibrium is an interest rate $R$, two densities $h^o(A, y)$ and $h^w(A, y)$, and two functions $c^o(A, y)$ and $c^w(A, y)$ such that $c^o(A, y)$ is the optimal policy of the objective agents and $c^w(A, y)$ is the optimal policy of the wishful thinkers, the goods market clears

$$\int_0^\eta c^o(A_{it-1}, y_{it}) di + \int_\eta^1 c^w(A_{it-1}, y_{it}) di = \int_0^1 y_{it} di$$

and $h^o(A, y)$ and $h^w(A, y)$ characterize the steady state distribution across states of the two types of agent respectively.

We have the following proposition (the proof is in the appendix).

**Propostion** The following hold:

1. An equilibrium exists.

2. The consumption function $c^w(A, y)$ and the value function $V^w(A, y)$ for the wishful thinkers are increasing in both their arguments.

3. Whenever $A' > \phi$, the Euler equation for the wishful thinkers is

$$u'(c^w(A, y)) = \beta R E_{(A,y)} \left[ \frac{\exp\{\theta V^w(A', y')\}}{E_{(A,y)} \left[\exp\{\theta V^w(A', y')\}\right]} u'(c^w(A', y')) \right].$$

   where $E_{(A,y)}$ is the objective expectation conditional on $(A, y)$.

4. Given $y$, $h^o(A, y)$ first order stochastically dominates $h^w(A, y)$.

5. Whenever the equilibrium is unique, $R$ is decreasing in $\eta$.

Given the assumptions of DARA utility and i.i.d. income, Schechtman and Escudero (1977) show that asset accumulation is bounded above. The first three results then follow from standard dynamic programming arguments. Given that high value states tend to be high consumption states and high consumption states are low marginal utility states, the Euler equation implies that wishful thinkers tend to place greater weight on low marginal utility states than do the objective agents. It follows that $E_{(A,y)} \frac{u'(c^w(A',y'))}{u'(c^w(A,y))} < E_{(A,y)} \frac{u'(c^o(A',y'))}{u'(c^o(A,y))}$ whenever $A' > \phi$, and that wishful thinkers consume more, save less, and accumulate less wealth. The lower wealth explains point 4. The precautionary savings motive of the objective agents tends to push down the real interest rate. Wishful thinkers' optimism tends to push the interest rate up. $R$ is therefore increasing in the proportion of wishful thinkers which explains point 5.[15]

Note it is important here that wishful thinking affects the only the expectation of labor income in this Huggett economy. The bond return is fixed. In a model with risky assets, wishful thinkers can be optimistic about asset returns. Depending on the relative importance of income and substitution effects, saving could rise. We will revisit this issue below when we consider asset bubbles.

Figure 3 depicts the determination of the equilibrium. The gross interest rate $R$ is on the vertical axis. The two curves depict the average saving of each type of agent: $A^w(R)$ for wishful thinkers and $A^o(R)$ for objective agents. When $R$ equals 0, $A^o(R)$ equals $\phi$. As $R$ approaches $1/\beta$, $A^o(R)$ converges to $\infty$ (Chamberlain and Wilson, 2000). Since the objective agents save more, $A^o(R)$ is always to the right of $A^w(R)$. The equilibrium is at the point that total saving in the economy is equal to zero, so that $\eta A^o(R) + (1 - \eta) A^w(R) = 0$.

[15] The qualifier "whenever the equilibrium is unique" is needed only because one cannot guarantee that the equilibrium in the Huggett model is unique and hence there may be equilibria in which the $A(R)$ curves in Figure 4 below "bend the wrong way."

[Insert Figure 3]

## 5.2 Information-based trade

In our Huggett economy, agents have access to the same objective information yet hold different beliefs. The fact that they "agree to disagree" suggests that our model might lie outside of the class of considered by Milgrom and Stokey (1982). In fact, relative to Milgrom and Stokey, our model relaxes the assumption that it is common knowledge that all agents are rational expected utility maximizers.[16]

We present an example of information-based trade. There are four states $\{\omega_1, \omega_2, \omega_3, \omega_4\}$ and two assets $a_1$ and $a_2$. There are two agents: one is objective and one is a wishful thinker. The following table presents $u(a, \omega)$. These are the same for both agents.

| | $a_1$ | $a_2$ |
|---|---|---|
| $\omega_1$ | 6 | 4 |
| $\omega_2$ | 2 | 4 |
| $\omega_3$ | 2 | 3 |
| $\omega_4$ | 2 | 1 |

Table 1: Payoffs to the two actions

We assume that initially all states are equally likely. Since both assets have the same expected payout, the objective agent is indifferent. The wishful thinker, however, places more weight on $\omega_1$ and prefers asset $a_1$.[17] Therefore assigning $a_1$ to the wishful thinker and $a_2$ to the objective agent is ex ante efficient.

Suppose now that the agents receive a signal that indicates whether the state is in the set $\{\omega_1, \omega_2\}$ or the set $\{\omega_3, \omega_4\}$. If the signal indicates $\{\omega_1, \omega_2\}$ there is no trade. The objective agent is still indifferent and the wishful thinker still prefers $a_1$. If, however, the signal indicates $\{\omega_3, \omega_4\}$. The objective agent remains indifferent, but the wishful thinker places greater weigh on $\omega_3$ and prefers to trade $a_1$ for $a_2$. The objective agent obliges.

Note that the signal could be the private information of the wishful thinker without affecting this outcome. The wishful thinker's willingness to trade would reveal that their

[16] We also implicity shut down any learning mechanism that would get rid of biased beliefs in the long run as agents gather enough information to reject their biased model in favor or the objective one.

[17] For example if $\theta = 1$, the expected value of $a_1$ is 4.67 and the expected value of $a_2$ is 3.50.

information set is $\{\omega_3, \omega_4\}$, but the agents would agree to disagree on the relative probabilities of $\omega_3$ and $\omega_4$ and trade would still take place. Similarly we could raise the payoff to $a_1$ in state $\omega_3$ and the payoff to $a_2$ in state $\omega_2$ by a small amount $\varepsilon$ without altering the ex ante efficiency of the allocation, and then the signal could be the private information of the objective agent and trade would take place in the case that the signal was $\{\omega_3, \omega_4\}$.

When the signal is private information, agents learn from trade. They understand the motivation of the other agent and back out the signal from the other agents desire to trade. This differentiates our framework from other models of trade with heterogeneous beliefs in which agents hold their beliefs dogmatically and do not learn from trade. Examples include Blume and Easley (1992), Geanakoplos (2003), Borovicka (2018), or Caballero and Simsek (2018).

While special this example illustrates the general point that wishful thinkers and objective agents process information differently. A change in circumstance may alter the relative probability of various states in the mind of the objective agent, but it has an additional effect on the perspective of the wishful thinker as it also alters the wishful thinker's desire to distort these probabilities. In this way, news is likely to differentially affect perceived asset returns, and these differential effects on perceived returns are likely to generate trade.

### 5.3 Bubbles

We saw in our Huggett economy that the presence of wishful thinkers can drive up the price of an asset, in this case a risk free bond, relative to that of an economy populated only by objective agents. Here we use this observation to flesh out a story of asset bubbles. In normal times, assets will be priced mainly by the objective agents as we saw in our Huggett economy that the wishful thinkers will tend to have less wealth. Every now and then, however, the wishful thinkers may take over and an asset price may rise above fundamental value, where fundamental value is taken to be the price in an economy dominated by objective agents.

Kindelberger (1978) describes the typical bubble as follows. The bubble begins with the introduction of a new asset, typically reflecting a new technology such as railroads, information technology, or cryptocurrency. A period of good news then increases interest in the new asset. In the second phase, interest evolves into euphoria, and prices rise above what would be expected by an objective observer. This is the bubble phase. Typically the bubble does not crash immediately. There is a period of hesitation at the top of the market. This is the third phase. The final phase is the crash, as reality sets in and prices return to normal.

A model with wishful thinkers and objective agents can fit this general pattern. Here we sketch the broad outlines of such a model. Consider the introduction of a new technology. Suppose that initially it is not known wheather the technology is viable or not, and that conditional on being viable there is a chance that the technology is transformative and a chance that the technology is merely mildly profitable. Wishful thinkers will be drawn to the new technology both because it is unproven and hence uncertain and because it promises high returns. Initially wishful thinkers will make up only a small portion of the demand for the asset because they tend to be less wealthy than objective agents. The initial success of the asset may cause their role to grow with time. This happens for two reasons. First, since the wishful thinkers will hold a greater share of the new asset in their portfolios, initial reports that the technology is viable will disproportionately benefit them.[18] Second, in the flow model each bit of good news will be interpreted in a positive light, leading to greater and greater optimism. As the importance of the wishful thinkers grows, they bid the price of the asset above fundamentals. As the price rises the importance of the wishful thinkers continues to grow, as objective agents see the market as overvalued. At the top of the market, when an objective observer would question the potential of the technology to be transformative, wishful thinkers will tend to downplay bad news. This is the period of hesitation. Wishful thinking, however, is not magical thinking. Eventually even optimists must admit that the asset is only mildly successful.

# 6 Discussion

## 6.1 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 problems (3) or (11), the agent foresees that their beliefs will change and takes this into consideration. Another possibility is that agents are naïve. 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 naïve. Optimism, overconfidence and polarization only depended on the choice of beliefs, not on

[18] This effect is also present in Caballero and Simsek (2018). The price of the risky asset rises with the wealth of their optimistic agents.

the choice of actions. Naïveté, however, gives rise to dynamic inconsistency, as the agent fails to anticipate how a choice today will affect beliefs that affect choices tomorrow. Time inconsistency gives rise to additional phenomenon which we consider below.

One such phenomenon is the endowment effect (Kahneman, Knetsch, and Thaler, 1990). The endowment effect is the idea that the mere possession of an object increases its value in the mind of the possessor. Consider a naïve agent in the cumulative model and consider an object of uncertain value. If the agent possesses the object, then the agent will tend to overweight the possibility that the object has high value. If the agent does not posses the object they have no reason to believe that it is of high value and since they are naïve they would not consider how possession of the object would affect their beliefs. This gives rise to the endowment effect. This explanation of the endowment effect is similar to those in cognitive science which emphasize biased search, memory and information processing (See Morewedge and Giblin (2015) for a review).[19]

Another phenomenon is the foot-in-the-door technique in marketing (Freedman and Fraser, 1966). 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 $\omega_H$ and $\omega_L$ and consider a gamble in which the agent gets $x$ if the state is $\omega_H$ and $-x$ otherwise. Suppose that utility, $u(x)$, is increasing and concave and that the agent has a prior belief $\mu_H$ that the state is $\omega_H$. If the agent is naïve, they base the decision on their prior and do not consider how taking the gamble will affect their beliefs. If $\mu_H$ is high enough, but not too high, it is possible due to the concavity of utility that the agent would choose the gamble for small $x$ and avoid the gamble for larger $x$. Upon choosing the gamble for small $x$, wishful thinking would lead the agent to increase the perceived probability of $\omega_H$, and it is then possible that they would be willing to take the larger gamble.

## 6.2 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 (3), but the cost (2) enters with the opposite sign and the agent first minimizes with respect

[19] Since the agent in the flow model manipulates the flow of information and not the stock of information, the flow model can only explain the endowment effect if the person who receives the object also receives a signal that they can manipulate.

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 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 gives rise to 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. While robustness appears to help explain asset pricing behavior, there are many economic situations where are better explained by wishful thinking. Corporate finance tends to treat entrepreneurs as optimistic. Referring to entrepreneurs, Daniel Kahneman said, “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. Hall and Woodward (2010) argue that due to the extreme dispersion in payoffs, an entrepreneur backed by venture capital with rational expectations and a coefficient of relative risk aversion equal to two should place a certainty equivalent value only slightly greater than zero on the distribution of outcomes that they face at the time that they start their company. 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.

Many self-control problems would appear to be more consistent with optimism than robustness. We have already cited the evidence on gym memberships (DellaVigna and Malmendier, 2006). Payday lending is another example. 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 also tend to be optimistic regarding how many times they will roll over debt.

It is a question for future research to find the key determinants of when a domain is more appropriate for wishful thinking and when a domain is more appropriate for robustness. Our Huggett economy suggests that in normal times asset pricing may naturally be a domain for robustness, since wishful thinkers tend to accumulate less wealth. Our discussion of bubbles suggests that sometimes, however, wishful thinkers may come to play a larger role. Entrepreneurship, on the other hand, would seem to be a natural domain for wishful thinking.

It may also be the case that the same agents are wishful thinkers in some situations and robust in others. Bassanin, Faia, and Valaria (2018) take a step in this direction. Citing psychological research that suggests people are sometimes ambiguity averse and sometimes ambiguity seeking and that these attitudes are state dependent, they construct a business cycle model that incorporates both behaviors. They assume that agents are ambiguity averse if the value function is below its historical mean, and ambiguity seeking otherwise. Their model amplifies business cycles by generating optimism in booms and pessimism in busts.

## 6.3 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. Like our agents in the flow model, their agents are subjective Bayesians. Their model leads to many of the same phenomenon as our model, in particular optimism and overconfidence.

The main difference between the two approaches is the way in which they handle the cost of belief choice. We assume a cost of distorting beliefs from some objective benchmark. Brunnermeier and Parker focus instead on how distorted beliefs might lead to suboptimal decisions. Because an agent with distorted beliefs may not be aware that their decisions are suboptimal, they assume that the costs and benefits of belief choice are evaluated according to the objective beliefs. To avoid the contractions arising from an agent using the objective beliefs for one set of decisions and the distorted beliefs for another set of decisions, they separate the two decisions in time. In their model there is an initial "period zero" that occurs before all choices are made. In this period the agent chooses a prior. This choice balances the utility gain from choosing an optimistic prior against the against the mistakes that result from a mistaken prior. This choice is made using the objective probabilities to weight outcomes. In subsequent periods, the agent observes the world, updates their information as would a Bayesian and makes decisions. There is no subsequent belief choice.

Suboptimal decisions are clearly an important cost of distorted beliefs. Incorporating this cost, however, presents significant modelling challenges as one has to figure out how to model an agent that simultaneously juggles two sets of beliefs. Brunnermeier and Parker found an ingenious solution to this dilemma by separating the choice of beliefs from the choice of actions. The cost of their approach is that beliefs are only chosen once at the beginning of life.

Other differences between the models evolve out of the placement of belief choice at the beginning of life. In Brunnermeier and Parker, agents have an incorrect prior, but their interpretation of evidence accords with objective reality. In our flow model, agents may or may not have an incorrect prior, it is there interpretation of signals that is overly optimistic. In Brunnermeier and Parker the initial choice of beliefs tends to affect future choices. In our model, past choices also mold future beliefs.

# 7 Conclusion

We model an agent who gets utility from their beliefs and therefore interprets information optimistically. The framework can explain behavioral biases such as optimism, procrastination, confirmation bias, polarization, the endowment effect, and the foot-in-the-door phenomenon. In spite of these biases, the agent is subjectively Bayesian in some formulations.

Our theory is based on two fundamental ideas. First that agents derive utility from their beliefs along the lines of Jevons (1905), Loewenstein (1987) and Caplin and Leahy (2001). The second is that at any point in time there are a set of models of the world that are all plausible (Hansen and Sargent, 2008), so that agents have some freedom in choosing their beliefs without choosing beliefs that are obviously wrong.

An interesting direction for future research is to endogenize the set of plausible models. This could be done either on the supply side or the demand side. If one takes the view that the set of plausible models is well represented by the views in the mainstream media, one could endogenize this supply along the lines Mullainathan and Schleifer (2005). On the demand side, one could imagine enriching the model to add a choice of attention along the lines of Sims (1998), Matejka and McKay (2015), or Caplin, Csaba, Leahy and Nov (2018). Since wishful thinkers have a preference for late resolution of uncertainty, one might expect wishful thinkers to exhibit willful ignorance.

## References


Akerlof, George, and William Dickens (1982), "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 72, 307-319.

Bassanin, Marzio, Ester Faia, and Valeria Patella (2018), "Ambiguous Leverage Cycles," working paper.

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.

Blume, Lawrence, and David Easley (1992), "Evolution and Market Behavior." *Journal of Economic Theory* 58, 9–40.

Brunnermeier, Markus and Jonathan Parker (2005), "Optimal Expectations," *American Economic Review*, 1092-1118.

Brunnermeier, Markus, Filippos Papakonstantinou, and Jonathan Parker (2008), "An Economic Model of the Planning Fallacy," NBER Working Paper 14228.

Burnside, Craig, Martin Eichenbaum and Sergio Rebelo (2016), "Understanding Housing Booms and Busts," *Journal of Political Economy* 124, 1088-1147.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Schleifer (2018), "Diagnostic Expectations and Credit Cycles," *Journal of Finance* 73, 199-227.

Borovicka, Jaroslav (2018), "Survival and Long-Run Dynamics with Heterogeneous Beliefs under Recursive Preferences," *Journal of Political Economy,* forthcoming.

Buehler, Roger, Dale Griffin, and Michael Ross (1994), "Exploring the Planning Fallacy: Why People Underestimate their Task Completion Times," *Journal of Personality and Social Psychology 67*, 366-381.

Buehler, Roger, Dale Griffin, and Heather MacDonald (1997), "The Role of Motivated Reasoning in Optimistic Time Predictions," *Personality and Social Psychology Bulletin* 23, 238-247.

Caballero, Ricardo, and Alp Simsek (2108), "A Risk-centric Model of Demand Recessions and Macroprudential Policy," NBER Working Paper No. 23614

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.

Caplin, Andrew, Daniel Csaba, John Leahy, and Oded Nov (2018), "Rational Inattention and Psychometrics," working paper.

Chamberlain, Gary and Charles Wilson (2000), "Optimal Intertemporal Consumption under Uncertainty." *Review of Economic Dynamics* 3, 365-395.

Cooper, Arnold, Carolyn Woo, and William Dunkelburg (1988), "Entrepreneur's Perceived Chances of Success," *Journal of Business Venturing* 3, 97-108.

Dixit, Avinash (1989), "Entry and Exit Decisions under Uncertainty," *Journal of Political Economy* 97, 620-638.

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.

DellaVigna Stefano, and Ulrike Malmendier (2006), "Paying not to go to the Gym," *American Economic Review* 96, 694-719.

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.

Exley, Christine, and Judd Kessler (2018), "Motivated Errors," working paper.

Freedman, J. L., and S. C. Fraser (1966), "Compliance without pressure: The foot-in-the-door technique." *Journal of Personality and Social Psychology* 4, 195–202.

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.

Geanakoplos, John (2003) "Liquidity, Default, and Crashes: Endogenous Contracts in General Equilibrium," in *Advances in Economics and Econometrics: Theory and Applications*, Econometric Society Monographs, Eighth World Conference 2, 170–205. New York: Cambridge University Press.

Hall, Robert, and Susan Woodward (2010), "The Burden of the Nondiversifiable Risk of Entrepreneurship," *American Economic Review* 100, 1163-1194.

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.

Hopenhayn, Hugo (1992), "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica* 60, 1127-1150.

Huggett, Mark (1993), "The Risk Free Rate in Heterogeneous-Agent, Incomplete-Insurance Economies," *Journal of Economic, Dynamics and Control* 17, 953-969.

Ito, Takatoshi (1990), "Foreign Exchange Rate Expectations: Micro Survey Data," *American Economic Review* 80, 434-449.

Jevons, William (1905), *Essays in Economics*, London: Macmillan.

Kahneman, Daniel, Jack Knetsch, and Richard Thaler (1990), "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy* 98, 1325–1348.

Kahneman, Daniel, and Amos Tversky (1979): "Intuitive Prediction: Biases and corrective procedures," *TIMS Studies in Management Science*, 12,313–327.

Kindelberger, Charles (1978), *Manias, Crashes and Panics,* New York: Basic Books.

Knox, R. E., Inkster, J. A. (1968), "Postdecision Dissonance at Posttime," *Journal of Personality and Social Psychology*, 18, 319–323.

Loomes, Graham and Robert Sugden (1982), "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *The Economic Journal* 92, 805-824.

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.

Schechtman, Jack, and Vera Escudero (1977), "Some Results on 'An Income Fluctuation Problem'," *Journal of Economic Theory* 16, 151-166.

Stokey, Nancy, and Robert Lucas (1989), *Recursive Methods in Economic Dynamics*, Cambridge: Harvard University Press.

Milgrom, Paul, and Nancy Stokey (1982), "Information, Trade, and Common Knowledge," *Journal of Economic Theory* 26, 17-27.

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.

Morewedge, Carey, and Colleen Giblin (2015), "Explanations of the endowment effect: an integrative review," *Trends in Cognitive Sciences* 19, 339–348.

Matejka, Filip, and Alisdair McKay (2015), "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105, 272-298.

Mullainathan, Sendil, and Andrei Schleifer (2005), "The Market for News," *American Economic Review* 95, 1031-1052.

Pronin, Emily, Kimberly McCarthy, Sylvia Rodriguez, and Daniel Wegner (2006), "Everyday Magical Powers: The Role of Apparent Mental Causation in the Overestimation of Personal Influence," *Journal of Personality and Social Psychology* 91, 218 –231

Sims, Christopher (1998), "Stickiness," *Carnegie-Rochester Conference Series on Public Policy* 49, :317-356.

# 8 Appendix

## 8.1 Derivation of (9)

Consider the maximization problem (8). The first order condition for $p(s|\omega)$ is

$$0 = \frac{\partial E_\gamma u(\omega)}{\partial p(s|\omega)} - \frac{1}{\theta} \ln \frac{p(s|\omega)}{\bar{p}(s|\omega)} - \frac{1}{\theta} - \lambda_\omega$$

where $\lambda$ is the Lagrange multiplier on the constraint $\sum_{s'} p(s'|\omega) = 1$. We ignore the constraint $p(\omega) \geq 0$. We will show that the it is never binding.

Solving for $p(s|\omega)$,

$$p(s|\omega) = \bar{p}(s|\omega) \exp\left[\theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} - 1 - \theta\lambda_\omega\right] \tag{12}$$

Now consider $s' \in S\backslash s$. The first order condition for $p(s'|\omega)$ implies

$$p(s'|\omega) = \bar{p}(s'|\omega) \exp\left[-1 - \theta\lambda_\omega\right]$$

as $\frac{\partial E_\gamma u(\omega)}{\partial p(s'|\omega)} = 0$.

Summing over all $s'$ including $s$,

$$1 = \sum_{s' \in S} p(s'|\omega) = \bar{p}(s|\omega) \exp\left[\theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} - 1 - \theta\lambda_\omega\right] + \sum_{s' \in S\backslash S} \bar{p}(s'|\omega) \exp\left[-1 - \theta\lambda_\omega\right] \tag{13}$$

or

$$\exp\left[-1 - \theta\lambda_\omega\right] = \left[\bar{p}(s|\omega) \exp\left[\theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right] + (1 - \bar{p}(s'|\omega))\right]^{-1}$$

Dividing the left-hand side of (12) by the left-hand side of (13) yields,

$$p(s|\omega) = \frac{\bar{p}(s|\omega) \exp\left[\theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]}{\bar{p}(s|\omega) \exp\left[\theta \frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right] + (1 - \bar{p}(s|\omega))}$$

as required. Note that given $\bar{p}(s|\omega) \in (0, 1)$, $p(s|\omega) > 0$ so the constraint $p(s|\omega) \geq 0$ is never binding.

## 8.2 Derivation of the Value of action $a$

Substituting (9) into (7)

$$
\begin{aligned}
& -\frac{1}{\theta}\sum_{\omega\in\Omega}\sum_{s'\in S} p(s'|\omega)\ln\frac{p(s'|\omega)}{\bar{p}(s'|\omega)} \\
= & -\frac{1}{\theta}\left[\sum_{\omega\in\Omega} p(s|\omega)\ln\frac{\frac{\bar{p}(s|\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]}{\bar{p}(s|\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]+(1-\bar{p}(s|\omega))}}{\bar{p}(s|\omega)} + \sum_{s'\in S\setminus s}\sum_{\omega\in\Omega} p(s'|\omega)\ln\frac{\frac{\bar{p}(s'|\omega)}{\bar{p}(s|\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]+(1-\bar{p}(s|\omega))}}{\bar{p}(s'|\omega)}\right] \\
= & -\sum_{\omega\in\Omega} p(s|\omega)\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} + \frac{1}{\theta}\left[\sum_{\omega\in\Omega} p(s|\omega)\ln\left(\bar{p}(s|\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right] + (1-\bar{p}(s|\omega))\right) + \sum_{s'\in S\setminus s}\sum_{\omega\in\Omega} p(s'|\omega \right. \\
= & -\sum_{\omega\in\Omega} p(s|\omega)\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} + \frac{1}{\theta}\sum_{\omega\in\Omega}\ln\left(\bar{p}(s|\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right] + (1-\bar{p}(s|\omega))\right)
\end{aligned}
$$

Focus on the first term

$$
\sum_{\omega\in\Omega} p(s|\omega)\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)} = \sum_{\omega\in\Omega} p(s|\omega)\left(\frac{\mu(\omega)}{\sum_{\omega'} p(s|\omega')\mu(\omega')}[u(\omega) - E_\gamma u(\omega)]\right) = 0
$$

which completes the derivation.

## 8.3 Derivation of (11)

We begin with the maximization problem

$$
\begin{aligned}
V(A,y) \quad = \quad & \max_{\{\gamma(y^t)\}_{t>0,y\in Y},\{A(y^t)\}_{t>0,y\in Y}} \sum_{t=0}^{\infty}\beta^t\sum_{y\in Y}\gamma(y^t)u(RA(y^{t-1}(y^t)) + y_t - A(y^t)) \\
& -\frac{1-\beta}{\theta}\sum_{t=1}^{\infty}\beta^t\sum_{y^t\in Y^t}\gamma(y^t)\ln\frac{\gamma(y^t)}{\bar{\gamma}(y^t)}
\end{aligned}
$$

where $y^{t-1}(y^t)$ denotes the history through $y^{t-1}$ embedded in $y^t$ and $A > \phi$. The key observation that allows us to write the problem recursively is that

$$\begin{aligned}\sum_{y^t \in Y^t} \gamma(y^t) \ln \frac{\gamma(y^t)}{\bar{\gamma}(y^t)} &= \sum_{y_1 \in Y} \left[ \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y_1)\gamma(y^t|y_1) \ln \left( \frac{\gamma(y_1)\gamma(y^t|y_1)}{\bar{\gamma}(y_1)\bar{\gamma}(y^t|y_1)} \right) \right] \\ &= \sum_{y_1 \in Y} \left[ \gamma(y_1) \ln \frac{\gamma(y_1)}{\bar{\gamma}(y_1)} + \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y_1)\gamma(y^t|y_1) \ln \frac{\gamma(y^t|y_1)}{\bar{\gamma}(y^t|y_1)} \right]\end{aligned}$$

so that

$$\begin{aligned}&\frac{1-\beta}{\theta} \sum_{t=0}^{\infty} \beta^t \sum_{y^t \in Y^t} \gamma(y^t) \ln \frac{\gamma(y^t)}{\bar{\gamma}(y^t)} \\ &= \frac{1-\beta}{\theta} \sum_{t=1}^{\infty} \beta^t \sum_{y_1 \in Y} \left[ \gamma(y_1) \ln \frac{\gamma(y_1)}{\bar{\gamma}(y_1)} + \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y_1)\gamma(y^t|y_1) \ln \frac{\gamma(y^t|y_1)}{\bar{\gamma}(y^t|y_1)} \right] \\ &= \frac{\beta}{\theta} \sum_{y_1 \in Y} \gamma(y_1) \ln \frac{\gamma(y_1)}{\bar{\gamma}(y_1)} + \sum_{y_1 \in Y} \beta\gamma(y_1) \frac{1-\beta}{\theta} \sum_{t=2}^{\infty} \beta^{t-1} \left[ \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y^t|y_1) \ln \frac{\gamma(y^t|y_1)}{\bar{\gamma}(y^t|y_1)} \right]\end{aligned}$$

Similarly we can rewrite expected utility

$$\begin{aligned}&\sum_{t=0}^{\infty} \beta^t \sum_{y \in Y} \gamma(y^t) u(RA(y^t) - A(y^{t-1}(y^t))) \\ &= u(RA(y_1) - A_0) + \sum_{t=1}^{\infty} \beta\beta^{t-1} \sum_{y_1 \in Y} \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y_1)\gamma(y^t|y_1) u(RA(y^t) - A(y^{t-1}(y^t))) \\ &= u(RA(y_1) - A_0) + \beta \sum_{y_1 \in Y} \gamma(y_1) \sum_{t=1}^{\infty} \beta^{t-1} \sum_{y^t \in Y^t \text{ s.t. } y_1 \in y^t} \gamma(y^t|y_1) u(RA(y^t) - A(y^{t-1}(y^t)))\end{aligned}$$

Putting the two together

$$
\begin{aligned}
V(A,y) &= \max_{\{\gamma(y^t)\}_{t>0,y\in Y},\{A(y^t)\}_{t>0,y\in Y}} \sum_{t=0}^{\infty}\beta^t\sum_{y\in Y}\gamma(y^t)u(RA(y^t)-A(y^{t-1}(y^t))) \\
&\qquad -\frac{1-\beta}{\theta}\sum_{t=1}^{\infty}\beta^t\sum_{y^t\in Y^t}\gamma(y^t)\ln\frac{\gamma(y^t)}{\bar{\gamma}(y^t)} \\
&= \max_{\{\gamma(y^t)\}_{t>0,y\in Y},\{A(y^t)\}_{t>0,y\in Y}} u(RA(y_1)-A_0)-\frac{\beta}{\theta}\sum_{y_1\in Y}\gamma(y_1)\ln\frac{\gamma(y_1)}{\bar{\gamma}(y_1)} \\
&\quad +\beta\sum_{y_1\in Y}\gamma(y_1)\left[\sum_{t=1}^{\infty}\beta^{t-1}\sum_{y^t\in Y^t \text{ s.t. } y_1\in y^t}\gamma(y^t|y_1)u(RA(y^t)-A(y^{t-1}(y^t)))\right. \\
&\quad \left.-\frac{1-\beta}{\theta}\sum_{t=2}^{\infty}\beta^{t-1}\sum_{y^t\in Y^t \text{ s.t. } y_1\in y^t}\gamma(y^t|y_1)\ln\frac{\gamma(y^t|y_1)}{\bar{\gamma}(y^t|y_1)}\right] \\
&= \max_{\{\gamma(y^t)\}_{t>0,y\in Y},\{A(y^t)\}_{t>0,y\in Y}} u(RA(y_1)-A_0) \\
&\qquad -\frac{\beta}{\theta}\sum_{y_1\in Y}\gamma(y_1)\ln\frac{\gamma(y_1)}{\bar{\gamma}(y_1)}+\beta\sum_{y_1\in Y}\gamma(y_1)V(A(y_0),y_1)
\end{aligned}
$$

which is the desired equation.

## 8.4 Proof of the Propositions

**Proposition 1** Given $\mu$ and $\bar{p}$, there exists a $p$ that satisfies (9) for all $\omega\in\Omega$.

**Proof:** Consider the mapping $T:[0,1]^{|\Omega|}\to[0,1]^{|\Omega|}$ defined by (9). Given any $p\in[0,1]^{|\Omega|}$ define $T(p)(\omega)\in[0,1]$, by

$$
T(p)(\omega)=\frac{\bar{p}(\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]}{\bar{p}(\omega)\exp\left[\theta\frac{\partial E_\gamma u(\omega)}{\partial p(\omega)}\right]+(1-\bar{p}(\omega))}.
$$

Clearly $T(p)(\omega)\in[0,1]$ for each $\omega$. 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. $p_H>\bar{p}_H$ and $p_L<\bar{p}_L$
2. $p_H$ is strictly increasing in $u_H-u_L$, whereas $p_L$ is strictly decreasing in $u_H-u_L$

3. $p_H$ is strictly increasing in $\theta$, whereas $p_L$ is strictly decreasing in $\theta$

4. $p_H$ is strictly increasing in $\bar{p}_H$, whereas $p_L$ is strictly increasing in $\bar{p}_L$

5. $p_H$ is strictly increasing in $\mu_H$ if $p_H\mu_H < p_L\mu_L$ and decreasing in $\mu_H$ if $p_H\mu_H > p_L\mu_L$. The opposite applies for $p_L$.

**Proof:** Given two states (9) can be rewritten as.

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

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

Consider now $T$ mapping defined in Proposition 1 and applied to the example with two states. Let $T(p_H)$ denote $T(p)(\omega_H)$. Then

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

Note that when $p_H = 0$, we have

$$T(0) = \frac{\bar{p}_H}{\bar{p}_H + (1-\bar{p}_H)e^{-\frac{\theta\mu_H(u_H-u_L)}{p_L\mu_L}}} > 0$$

and when $p_H = 1$

$$T(1) = \frac{\bar{p}_H}{\bar{p}_H + (1-\bar{p}_H)e^{-\frac{\theta\mu_H\mu_L p_L(u_H-u_L)}{(\mu_H+p_L\mu_L)^2}}} < 1$$

Each fixed point of the $T$ mapping is a solution to the first order condition.

Now consider the value of the policy $p_H$:

$$\frac{p_H\mu_H u_H}{p_H\mu_H + p_L\mu_L} - \frac{1}{\theta}p_H \ln\left(p_H/\bar{p}_H\right) - \frac{1}{\theta}(1-p_H)\ln\left(\left(1-p_H\right)/\left(1-\bar{p}_H\right)\right) + \text{terms independent of } p_H$$

The derivative of this value is

$$\begin{aligned} & \frac{\mu_H u_H}{p_H\mu_H + p_L\mu_L} - \frac{p_H\mu_H u_H + p_L\mu_L u_L}{(p_H\mu_H + p_L\mu_L)^2}\mu_H - \frac{1}{\theta}\ln\left(p_H/\bar{p}_H\right) + \frac{1}{\theta}\ln\left(\left(1-p_H\right)/\left(1-\bar{p}_H\right)\right) \\ = & \frac{\mu_H\mu_L p_L(u_H-u_L)}{(p_H\mu_H + p_L\mu_L)^2} - \frac{1}{\theta}\ln\left(p_H/\bar{p}_H\right) + \frac{1}{\theta}\ln\left(\left(1-p_H\right)/\left(1-\bar{p}_H\right)\right) \end{aligned}$$

Note here that $p_H$ and $p_L$ are chosen separately and do not necessarily sum to one. This derivative is positive and infinite if $p_H = 0$ and negative and infinite when $p_H = 1$. 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 $\mu_H - \mu_L$, $\theta$, and $\bar{p}_H$ together with the observation that $T'(p) < 1$ at all maxima.

As concerns point (5), consider the effect of an increase in $\mu_H$ 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

$$\begin{aligned}
(1-2\mu_H) \quad & (p_H \mu_H + p_L \mu_L) - 2(\mu_H - \mu_H^2)(p_H - p_L) \\
= \quad & p_H \mu_H + p_L \mu_L - 2\mu_H p_H \mu_H - 2\mu_H p_L \mu_L - 2\mu_H p_H + 2\mu_H p_L - 2\mu_H^2 p_L + 2\mu_H^2 p_H \\
= \quad & \mu_L p_L - \mu_H p_H
\end{aligned}$$

So $\frac{\mu_H(1-\mu_H)}{(p_H \mu_H + p_L(1-\mu_H))^2}$ is increasing in $\mu_H$ when $\mu_H p_H < \mu_L p_L$ and decreasing otherwise. It follows that $T(p)$ is increasing in $\mu_H$ when $\mu_H p_H < \mu_L p_L$ and decreasing otherwise, which implies (5).■

**Proposition 3** The following hold

1. An equilibrium exists.

2. The consumption function $c^w(A, y)$ and the value function $V^w(A, y)$ for the wishful thinkers are increasing in both their arguments.

3. Whenever $A' > \phi$, the Euler equation for the wishful thinkers is

$$u'(c^w(A, y)) = \beta R E_{(A,y)} \left[ \frac{\exp\{\theta V^w(A', y')\}}{E_{(A,y)} \left[ \exp\{\theta V^w(A', y')\} \right]} u'(c^w(A', y')) \right].$$

   where $E_y$ is the objective expectation conditional on $y$ and $V^w$ is the value of an optimal policy.

4. Given $y$, $h^o(A, y)$ first order stochastically dominates $h^w(A, y)$.

5. $R$ is decreasing in $\eta$.

**Proof:** We begin with the individual's problem and then prove existence of an equilibrium. Suppose that there an upper bound $\bar{A}$ such that $A_{it} \leq \bar{A}$. We will establish that such an $\bar{A}$ exists. Let $\Omega = [\phi, \bar{A}]$. Since $\phi$ is above the natural borrowing constraint by assumption, an individual with assets $\phi$ and the lowest income realization can always consume. This ensures that the utility function is bounded on $\Omega$.

Because income is i.i.d. the agent's current state is simply $a = RA + y$. The Bellman equation for the objective individuals is

$$V^o(a) = \max_{A'>\phi} u(a - A') + \beta E_y V^o(RA' + y')$$

Consider the mapping $TV^o(a) = \max_{A'>\phi} u(a - A') + \beta E_y\{V^o(RA' + y')\}$. Given $R < 1/\beta$, $u' > 0$, $u'' < 0$, standard dynamic programming arguments (Stokey and Lucas, 1979) insure that $T$ has a unique fixed point and that this is the optimal policy of the objective individual. Given the strict concavity of utility, the Theorem of the Maximum the maximum implies $A^{o\prime}(a|R)$ is continuous in $R$. Standard dynamic programming arguments also imply that $V^o(a)$ is increasing and differentiable in $a$.[20]

These results also apply to the wishful thinkers. The Bellman equation for the wishful thinkers is

$$V^w(a) = \max_{A'>\phi} u(a - A') + \beta \ln E \exp\{V^w(RA' + y')\}$$

Consider the mapping $TV^w(a) = \max_{A'>\phi} u(a - A') + \beta \ln E \exp\{V^w(RA' + y')\}$. Note that

$$\max_{A'\in\phi} u(a-A')+\beta \ln E \exp\{V^w(RA'+y')+a\} = \max_{A'\in\phi} u(a-A')+\beta(E \exp\{V^w(RA'+y')\})+\beta a$$

so that Blackwell's conditions hold (monotonicity is obvious), and $T$ is a contraction mapping. Again standard dynamic programming arguments insure that a solution exists $V^w(A, y)$, that $V^w(a)$ is increasing and differentiable in $a$ and that $A^{w\prime}(a|R)$ is continuous in $R$. This establishes point (2).

The Euler equation for the objective agents is standard

$$u'(c_t) = \beta R E u'(c_{t+1})$$

[20] $T^o$ maps the set of bounded weakly increasing functions into the set of bounded weakly increasing functions. The Benviniste-Schenkman theorem establishes differentiability.

The Euler equation for the wishful thinkers is less standard. The first order condition for $A'$ is

$$u'(c_t) = \beta R \frac{E \exp\{V^w(a_{t+1})\} \frac{d}{da} V^w(a_{t+1})}{E \exp\{V^w(a_{t+1})\}}$$

The derivative of the value function with respect to $A$ is

$$\frac{d}{da} V^w(a_t) = u'(c_t)$$

Together these imply.

$$u'(c_t) = \beta R \frac{E \exp\{V^w(a_{t+1})\} u'(c_{t+1})}{E \exp\{V^w(a_{t+1})\}}$$

This establishes point (3).

We now establish that $A^{w\prime}(a|R) \leq A^{o\prime}(a|R)$. Consider the mappings $T^o$ and $T^w$. Suppose that $\frac{d}{da} V^w(a_{t+1}) \leq \frac{d}{da} V^o(a_{t+1})$. $\frac{d}{da} TV^w(a_t) \leq \frac{d}{da} TV^o(a_t)$. Consider the Euler Equation of the wishful thinker

$$\frac{d}{da} TV^w(a_t) = \beta R \frac{E \exp\{V^w(a_{t+1})\} \frac{d}{da} V^w(a_{t+1})}{E \exp\{V^w(a_{t+1})\}}$$

We can expand the right-hand side

$$\begin{aligned}
\beta R \frac{E \exp\{V^w(a_{t+1})\} \frac{d}{da} V^w(a_{t+1})}{E \exp\{V^w(a_{t+1})\}} &= \beta R \frac{E_{y_t} \exp\{V^w(a_{t+1})\} E \frac{d}{da} V^w(a_{t+1}) + \operatorname{cov}\left(V^w(a_{t+1}), \frac{d}{da} V^w(a_{t+1})\right)}{E \exp\{V^w(a_{t+1})\}} \\
&= \beta R \left( E \frac{d}{da} V^w(a_{t+1}) + \frac{\operatorname{cov}\left(V^w(a_{t+1}), \frac{d}{da} V^w(a_{t+1})\right)}{E \exp\{V^w(a_{t+1})\}} \right)
\end{aligned}$$

Since $V^w(a)$ and $c$ are both increasing in $y$ and the exponential is positive, the second covariance term is non-positive. It follows that

$$\frac{d}{da} TV^w(a_t) \leq \beta R E_{y_t} \frac{d}{da} V^w(a_{t+1}) \leq \beta R E_{y_t} \frac{d}{da} V^o(a_{t+1}) = TV^o(a_t)$$

It then follows from Corollary 3.1 in Stokey and Lucas that $\frac{d}{da} V^w(a) \leq \frac{d}{da} V^o(a)$ and it follows from the Euler equation that $A^{w\prime}(a|R) \leq A^{o\prime}(a|R)$. This establishes point (4).

We now turn to the question of existence. Given $R < 1/\beta$ and DARA utility there exists an upper bound on the assets of the objective agents (Schechtman and Escudero, 1977). Let $\bar{A}$ denote the least upper bound. Given $R$, we then have $a \in X = [R\phi + y_1, R\bar{A} + y_S]$. We consider the measurable space $(X, \mathcal{S})$ where $\mathcal{S}$ is the Borel $\sigma$-algebra on $X$. Let $P^o(\mathcal{A}|a)$ and $P^w(\mathcal{A}|a)$ denote the probability of transiting to the set $\mathcal{A} \in \mathcal{S}$ conditional on the state

$a \in X$:

$$P^o(\mathcal{A}|A) = \sum_{y' \in Y} p(y') I_{A^{o\prime}(a)+y' \in \mathcal{A}}.$$

Stokey and Lucas (1999) Theorem 9.14 implies that $P^o$ has the Feller property. A similar argument establishes that $P^w$ has the Feller property. Since $A^{o\prime}(a|R)$ and $A^{w\prime}(a|R)$ are increasing in $a$, it follows that $P^w$ and $P^o$ are both monotone. Finally note that an agent with $A = \phi$ who receives a sufficiently long string of the highest income realizations will eventually accumulate enough wealth to place themselves in the neighborhood of $\bar{A}$, and that an agent with $A = \bar{A}$ who receives a sufficiently long strong of the lowest income realizations will eventually accumulate enough wealth to place themselves in the neighborhood of $\phi$. Assumption 12.1 of Stokey and Lucas is therefore satisfied. Stokey and Lucas' Theorem 12.12 then establishes the existence of a unique invariant distributions $h^o(a|R)$ and $h^w(a|R)$.

Let $A^{\omega\prime}(a|R)$ and $A^{o\prime}(a|R)$ denote the optimal saving policies of the wishful thinkers and the objective agents respectively. Define

$$A^w(R) = \int A^{\omega\prime}(a|R) h^w(da|R)$$

and

$$A^o(R) = \int A^{o\prime}(a|R) h^o(da|R)$$

Consider $R \to \bar{R}$, Lucas and Stokey Theorem 12.13 states that $h^w(a|R)$ converges weakly to $h^w(a|\bar{R})$ and $h^o(A|R)$ converges weakly to $h^o(A|\bar{R})$. As $A^{\omega\prime}(a|R)$ and $A^{o\prime}(a|R)$ are continuous, $A^w(R) \to A^w(\bar{R})$ and $A^o(R) \to A^o(\bar{R})$ establishing the continuity of $A^o(R)$ and $A^w(R)$

Now if $R = 0$, then $A^w(R) = A^o(R) = \phi < 0$, so that $A^w(R) + A^o(R) < 0$. And as $R \to 1/\beta$, $A^o(R) \to \infty$ (Chamberlain and Wilson, 2000) Since $A^w(R) \geq \phi$, we have $A^w(R) + A^o(R) > 0$ for $R$ sufficiently close to $1/\beta$. Since $A^w(R)$ and $A^o(R)$ are continuous, there exists $R$ such that

$$A^w(R) + A^o(R) = 0.$$

This establishes point (1).

Consider $R$ as a function of $\eta$, the fraction of objective agents. An increase in $\eta$, given $R$, drives up consumption and down the interest rate. This establishes point (5).$\blacksquare$

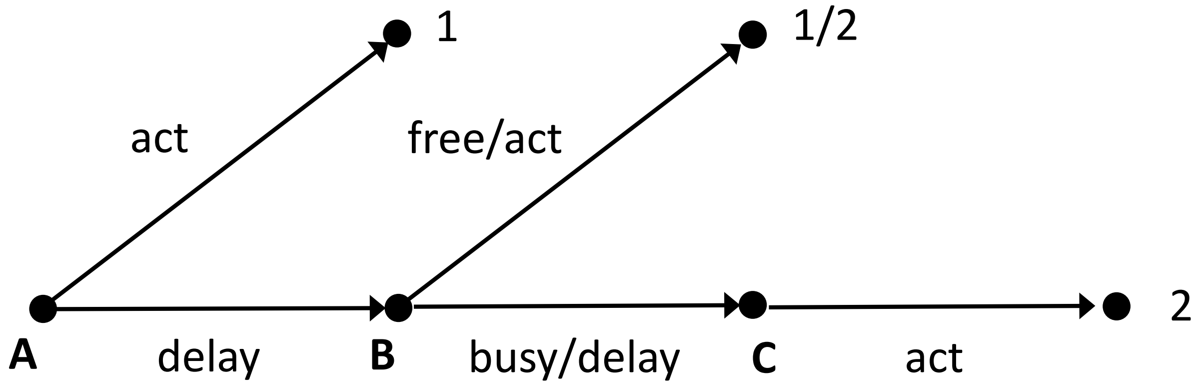


Figure 2: An Example of Procrastination

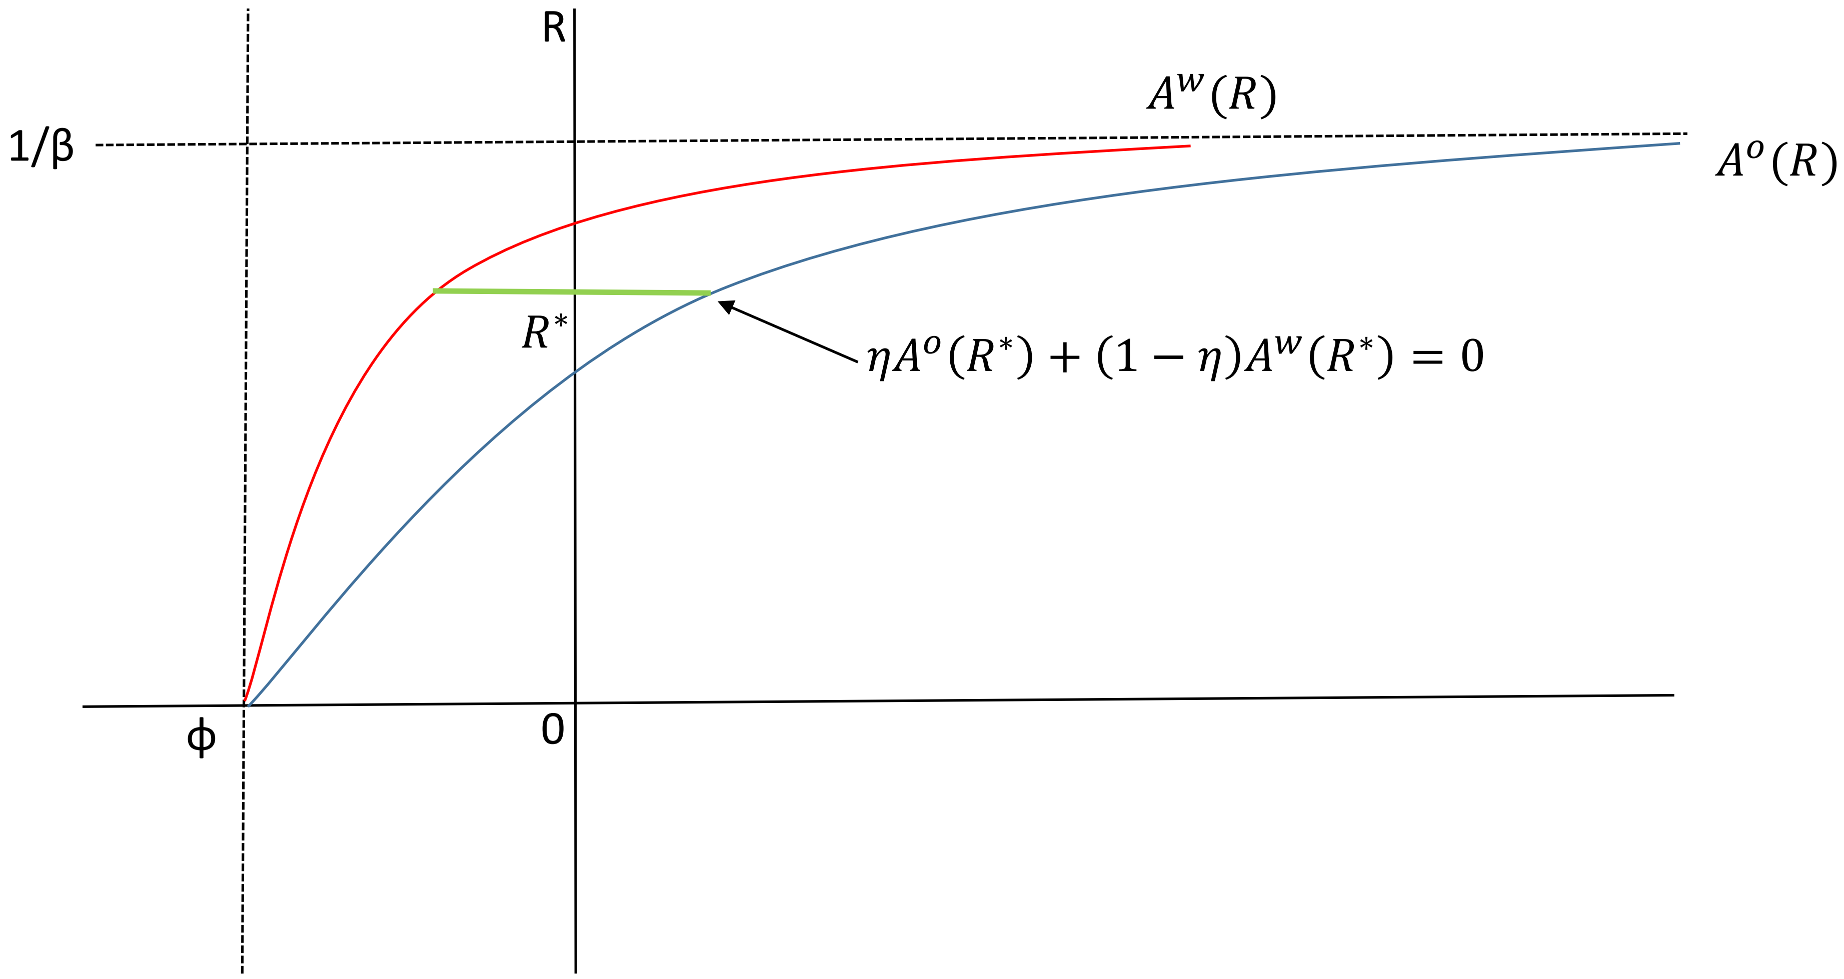


Figure 3: Equilibrium determination in the Huggett model