

NBER WORKING PAPER SERIES

OPTIMAL EXPECTATIONS

Markus K. Brunnermeier
Jonathan A. Parker

Working Paper 10707
<http://www.nber.org/papers/w10707>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2004

For helpful comments, we thank Andrew Abel, Roland Bénabou, Douglas Bernheim, Andrew Caplin, Larry Epstein, Ana Fernandes, Christian Gollier, Lars Hansen, David Laibson, Augustin Landier, Erzo Luttmer, Sendhil Mullainathan, Filippos Papakonstantinou, Wolfgang Pesendorfer, Larry Samuelson, Laura Veldkamp and three anonymous referees, as well as seminar participants at University of Amsterdam, UC Berkeley, Birkbeck, University of Bonn, Boston University, Carnegie-Mellon, University of Chicago, Columbia, the Federal Reserve Board, Harvard, the Institute for Advanced Study, London Business School, London School of Economics, MIT, University of Munich, New York University, University of Pennsylvania, Princeton, Rochester, Stanford, Tilburg University, the University of Wisconsin, Yale, and conference participants at the 2003 Summer Meetings of the Econometric Society, the NBER Behavioral Finance Conference April 2003, the 2003 Western Finance Association Annual Meetings, the NBER Economic Fluctuations and Growth meeting July 2003, the Minnesota Workshop in Macroeconomic Theory 2003, Bank of Portugal Conference on Monetary Economics 2004, CEPR Summer Symposium at Gerzensee 2004. Both authors acknowledge financial support from the National Science Foundation (SES 021-4445 and SES 009-6076). Parker also thanks the Sloan Foundation and an NBER Aging and Health Economics Fellowship through the National Institute on Aging (T32 AG00186). The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

©2004 by Markus K. Brunnermeier and Jonathan A. Parker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Expectations
Markus K. Brunnermeier and Jonathan A. Parker
NBER Working Paper No. 10707
August 2004
JEL No. D1, D8, E21, G11, G12

ABSTRACT

This paper introduces a tractable, structural model of subjective beliefs. Forward-looking agents care about expected future utility flows, and hence have higher current felicity if they believe that better outcomes are more likely. On the other hand, biased expectations lead to poorer decisions and worse realized outcomes on average. Optimal expectations balance these forces by maximizing average felicity. A small bias in beliefs typically leads to first-order gains due to increased anticipatory utility and only to second-order costs due to distorted behavior. We show that in a portfolio choice problem, agents overestimate the return on their investment and exhibit a preference for skewness. In general equilibrium, agents' prior beliefs are endogenously heterogeneous. Finally, in a consumption-saving problem with stochastic income, agents are both overconfident and overoptimistic.

Markus K. Brunnermeier
Department of Economics
Bendheim Center for Finance
Princeton University
Princeton, NJ 08544
markus@princeton.edu

Jonathan A. Parker
Department of Economics
Bendheim Center for Finance
Princeton University
Princeton, NJ 08544
and NBER
jparker@princeton.edu

1 Introduction

Modern psychology views human behavior as a complex interaction of cognitive and emotional responses to external stimuli that sometimes results in dysfunctional outcomes. Modern economics takes a relatively simple view of human behavior as governed by unlimited cognitive ability applied to a small number of concrete goals and unencumbered by emotion. The central models of economics allow coherent analysis of behavior and economic policy, but eliminate “dysfunctional” outcomes, and in particular the possibility that individuals might persistently err in attaining their goals. One area in which there is substantial evidence that individuals do consistently err is in the assessment of probabilities. In particular, agents often overestimate the probability of good outcomes, such as their success (Alpert and Raiffa (1982), Weinstein (1980), and Buehler, Griffin, and Ross (1994)).

We provide a structural model of subjective beliefs in which agents hold incorrect but optimal beliefs. These optimal beliefs differ from objective beliefs in ways that match many of the claims in the psychology literature about “irrational” behavior. Further, in the canonical economic models that we study, these beliefs lead to economic behaviors that match observed outcomes that have puzzled the economics literature based on rational behavior and common priors. Our approach has three main elements.

First, at any instant people care about current utility flow and expected future utility flows. While it is standard that agents that care about expected future utility plan for the future, forward-looking agents have higher current felicity if they are optimistic about the future. Agents that care about expected future utility flows are happier if they overestimate the probability that their investments pay off well or their future labor income is high.

The second crucial element of our model is that such optimism affects decisions and worsens outcomes. Distorted beliefs distort actions. For example, an agent cannot derive utility from optimistically believing that she will be rich tomorrow, while basing her

consumption-saving decision on rational beliefs about future income.

How are these forces balanced? We assume that subjective beliefs maximize the agent's expected well-being, defined as the time-average of expected felicity over all periods. This third key element leads to a balance between the first two – the benefits of optimism and the costs of basing actions on distorted expectations.

We illustrate our theory of optimal expectations using three examples. In general, a small bias in beliefs typically leads to first-order gains due to increased anticipatory utility and only to second-order costs due to distorted behavior. Thus, beliefs tend towards optimism – states with greater utility flows are perceived as more likely. Further, optimal expectations are less rational when biases have little cost in realized outcomes and when biases have large benefits in terms of expected future happiness.

More specifically, in a portfolio choice problem, agents overestimate the return of their investment and prefer skewed returns. Second, in general equilibrium, agents' prior beliefs are endogenously heterogeneous and agents gamble against each other. We show in an example that the expected return on the risky asset is higher than in an economy populated by agents with rational beliefs if the return is negatively skewed. Third, in a consumption-saving problem with quadratic utility and stochastic income, agents are overconfident and overoptimistic; early in life they consume more than implied by rational beliefs. In addition, Brunnermeier and Parker (2002) shows in a different economic setting that agents with optimal expectations can exhibit intertemporal preference reversal, a greater readiness to accept commitment, regret, and a context effect in which non-chosen actions can affect utility.

Psychological theories provide many channels through which the human mind is able to hold beliefs inconsistent with the rational processing of objective data. First, to the extent that people are more likely to remember better outcomes, they will perceive them as more likely in the future, leading to optimistic biases in beliefs as in our optimal expectations

framework.¹ Second, most human behavior is not based on conscious cognition but is automatic, processed only in the limbic system and not the cortex (Bargh and Chartrand (1999)). If automatic processing is optimistic, then the agent may naturally approach problems with optimistic biases. However, the agent may also choose to apply cognition to discipline belief biases when the stakes are large, as in our optimal expectations framework.

Our model of beliefs differs markedly from treatments of risk in economics. While early models in macroeconomics specify beliefs exogenously as naive, myopic, or partially updated (e.g. Nerlove (1958)), since Muth (1960, 1961) and Lucas (1976) nearly all research has proceeded under the rational expectations assumption that subjective and objective beliefs coincide. There are two main arguments for this. First, the alternatives to rationality lack discipline. But our model provides precisely this discipline for subjective beliefs by specifying an objective for beliefs, that they maximize well-being. The second argument is that agents have the incentive to hold rational beliefs (or act as if they do) because these expectations make agents as well off as they can be. However, this rationale for rational expectations relies upon agents caring about the future but at the same time having their expectations about the future not affect their current felicity, which we see as inconsistent. Our approach of optimal expectations takes into account the fact that agents care in the present about utility flows that are expected in the future in defining what beliefs are optimal.

Most microeconomic models assume that agents share common prior beliefs. This “Harsanyi doctrine” is weaker than the assumption of rational expectations that all agents’ prior beliefs are equal to the objective probabilities governing equilibrium dynamics. But like rational expectations, the common priors assumption is quite restrictive and does not allow agents to “agree to disagree” (Aumann (1976)). Savage (1954) provides axiomatic

¹In Mullainathan (2002) individuals have imperfect recall and form expectations as if they did not. In Piccione and Rubinstein (1997) individuals understand that they have imperfect recall, and in Bernheim and Thomadsen (2003) individuals additionally can influence the memory process to increase anticipatory utility.

foundations for a more general theory in which agents hold arbitrary prior beliefs, so agents can agree to disagree. But if beliefs can be arbitrary, theory provides little structure or predictive power. Optimal expectations provides discipline to the study of subjective beliefs and heterogeneous priors. Framed in this way, optimal expectations is a theory of prior beliefs for Bayesian rational agents.

The key assumption that agents derive current felicity from expectations of future pleasures has its roots in the origins of utilitarianism. Detailed expositions on anticipatory utility can be found in the work of Bentham, Hume, Böhm-Barwerk and other early economists. More recently, the temporal elements of the utility concept have re-emerged in research at the juncture of psychology and economics (Loewenstein (1987), Kahneman, Wakker, and Sarin (1997), Kahneman (2000)), and have been incorporated formally into economic models in the form of belief-dependent utility by Geanakoplos, Pearce, and Stacchetti (1989), Caplin and Leahy (2001), and Yariv (2001).² In particular, Caplin and Leahy (2000) shows that competitive equilibria are generically intertemporally sub-optimal and so opens the door for belief distortion to increase well-being.

Several papers in economics study related models in which forward-looking agents distort beliefs. In particular, Akerlof and Dickens (1982) models agents as choosing beliefs to minimize their discomfort from fear of bad outcomes. In a two-period model, agents with rational beliefs choose an industry to work in, understanding that in the second period they will distort their beliefs about the hazards of their work and perhaps not invest in safety technology. Second, Landier (2000) studies a two-period game in which agents choose a prior before receiving a signal and subsequently taking an action based on their updated beliefs. Unlike our approach, belief dynamics are not Bayesian; common to our approach, agents tend to save less and be optimistic about portfolio returns.³ Third, time

²Caplin and Leahy (2004) and Eliaz and Spiegel (2003) show that the forward-looking nature of utility raises problems for the revealed preference approach to behavior and the expected utility framework in the context of the acquisition of information.

³Similarly, in Eyster (2002), Rabin and Schrag (1999), and Yariv (2002) agents distort beliefs to be consistent with past choices or beliefs.

inconsistent preferences can make it optimal to strategically ignore information (Carrillo and Mariotti (2000)) or distort beliefs (Bénabou and Tirole (2002, 2004)). In the latter, unlike in our model, multiple selves play intra-personal games with imperfect recall, and actions serve as signals to future selves. Similarly, concerns about self-reputations also play a central role in Harbaugh (2002). Finally, there is a large literature on bounded rationality and incomplete memory. Some of these models suggest mechanisms for how individuals achieve optimal expectations in the face of possibly contradictory data.

The structure of the paper is as follows. In Section 2 we introduce and discuss the general optimal expectations framework. Subsequently, in Sections 3 through 5 we use the optimal expectations framework to study behavior in three different canonical economic settings. Section 3 studies a two-period two-asset portfolio choice problem and shows that agents hold beliefs that are biased towards the belief that their investments will pay off well. Section 4 shows that in a two-agent economy of this type with no aggregate risk, optimal expectations are heterogeneous and agents gamble against one another. Section 5 analyzes the consumption-saving problem of an agent with quadratic utility receiving stochastic labor income over time, and shows that the agent is biased towards optimism and is overconfident, and so saves less than a rational agent. Section 6 concludes. An appendix contains proofs of all propositions.

2 The optimal expectations framework

We choose to maintain many of the assumptions of canonical economic theory: agents optimize knowing the correct mapping from actions to payoffs in different states of the world. But we allow agents' assessments of probabilities of different states to depart from the objective probabilities.

This section defines our framework in two steps. First, we describe the problem of the agent given an arbitrary set of beliefs. At any point in time agents maximize felicity, the present discounted value of expected flow utilities. Second, we define optimal expectations

as the set of beliefs that maximize well-being in the initial period. Well-being is the expected time-average of the agent's felicity, and so is a function of the agent's beliefs and the actions these beliefs induce.

2.1 Optimization given beliefs

Consider a canonical class of optimization problems. In each period from 1 to T , agents take their beliefs as given and choose control variables, c_t , and the implied evolution of state variables, x_t , to maximize their happiness. We consider a world where the uncertainty can be described by a finite number of states, S .⁴ Let $\pi(s_t|\underline{s}_{t-1})$ denote the true probability that state $s_t \in S$ is realized after state history $\underline{s}_{t-1} := (s_1, s_2, \dots, s_{t-1}) \in \underline{S}_{t-1}$. We depart from the canonical model in that agents are endowed with subjective probabilities that may not coincide with objective probabilities. Conditional and unconditional subjective probabilities are denoted by $\hat{\pi}(s_t|\underline{s}_{t-1})$ and $\hat{\pi}(\underline{s}_t)$ respectively, and satisfy the basic properties of probabilities (precisely specified subsequently).

At time t , the agent chooses control variables, c_t , to maximize his felicity, given by

$$\hat{E}[U(c_1, c_2, \dots, c_T) | \underline{s}_t], \quad (1)$$

where $U(\cdot)$ is increasing and strictly quasi-concave and \hat{E}_t is the subjective expectations operator associated with $\{\hat{\pi}\}$ and given information available at t . The agent maximizes subject to a resource constraint

$$x_{t+1} = g(x_t, c_t, s_{t+1}), \quad (2)$$

$$h(x_{T+1}) \geq 0 \text{ and } x_0 \text{ is given,} \quad (3)$$

where $g(\cdot)$ gives the evolution of the state variable and is continuous and differentiable in x and c , and $h(\cdot)$ gives the endpoint condition. Denote the optimal choice of the control as $c^*(\underline{s}_t, \{\hat{\pi}\})$ and induced state variables as $x^*(\underline{s}_t, \{\hat{\pi}\})$.

⁴Appendix A defines optimal expectations for the situation with a continuous state space.

While the agent's problem is standard and general, we employ the specific interpretation that $\hat{E}[U(\cdot)|\underline{s}_t]$ is the felicity of the agent at time t . The felicity of the agent depends on expected future utility flows, or 'anticipatory' consumption, so that subjective conditional beliefs directly impact felicity. To clarify this point, consider the canonical model with time-separable utility flows and exponential discounting. In this case, felicity at time t ,

$$\hat{E}[U(c_{t-1}, c_t, \dots, c_T)|\underline{s}_t] = \beta^{t-1} \left(\sum_{\tau=1}^{t-1} \beta^{-\tau} u(c_{t-\tau}) + u(c_t) + \hat{E} \left[\sum_{\tau=1}^{T-t} \beta^{\tau} u(c_{t+\tau}) | \underline{s}_t \right] \right),$$

is the sum of memory utility from past consumption, flow utility from current consumption, and anticipatory utility from future consumption.

2.2 Optimal beliefs

Subjective beliefs are a complete set of conditional probabilities after any history of the event tree, $\{\hat{\pi}(s_t|\underline{s}_{t-1})\}$. We require that subjective probabilities satisfy four properties.

Assumption 1 (Restrictions on probabilities)

- (i) $\sum_{s_t \in S} \hat{\pi}(s_t|\underline{s}_{t-1}) = 1$
- (ii) $\hat{\pi}(s_t|\underline{s}_{t-1}) \geq 0$
- (iii) $\hat{\pi}(\underline{s}'_t) = \hat{\pi}(s'_t|\underline{s}'_{t-1}) \hat{\pi}(s'_{t-1}|\underline{s}'_{t-2}) \cdots \hat{\pi}(s'_1)$
- (iv) $\hat{\pi}(s_t|\underline{s}_{t-1}) = 0$ if $\pi(s_t|\underline{s}_{t-1}) = 0$.

Assumption 1(i) is simply that probabilities sum to one. Assumptions 1(i) – (iii) imply that the law of iterated expectations holds for subjective probabilities. Assumption 1(iv) implies that in order to believe that something is possible, it must be possible. That is, agents understand the underlying model and only misperceive the probabilities. For example, consider an agent choosing to buy a lottery ticket. The states of the world are the possible numbers of the winning ticket. An agent can believe that a given number will win the lottery. But the agent cannot believe in the nonexistent state that she will win the lottery if she does not hold a lottery ticket or even if there is no lottery. Note

that it is possible for the agent to believe that a possible event is impossible. But since we specify subjective beliefs conditional on all objectively possible histories, as in the axiomatic framework of Myerson (1986), the agent's problem is always well-defined.

We further consider the class of problems for which a solution exists and provides finite felicity for all possible subjective beliefs.

Assumption 2 (Conditions on agent's problem)

$$\hat{E}[U(c_1^*, c_2^*, \dots, c_T^*) | \underline{s}_t] < \infty \text{ for all } \underline{s}_t \text{ and for all } \{\hat{\pi}\} \text{ satisfying Assumption 1.}$$

Optimal expectations are the subjective probabilities that maximize the agent's life-time happiness. Formally, optimal expectations maximize well-being, \mathcal{W} , defined as the expected time-average of the felicity of the agent.

Definition 1 *Optimal expectations (OE) are a set of subjective probabilities $\{\hat{\pi}^{OE}(s_t | \underline{s}_{t-1})\}$ that maximize well-being*

$$\mathcal{W} := E \left[\frac{1}{T} \sum_{t=1}^T \hat{E}[U(c_1^*, c_2^*, \dots, c_T^*) | \underline{s}_t] \right] \quad (4)$$

subject to the four restrictions on subjective probabilities (Assumption 1).

In addition to being both simple and natural, this objective function is similar to what Caplin and Leahy (2000) argue should be the welfare function. Further, this choice of \mathcal{W} has the feature that under rational expectations well-being coincides with the agent's felicity, so the agent's actions maximize both well-being and felicity. We further discuss these issues in Subsection 2.3.

Optimal expectations exist if $c^{OE}(\underline{s}_t)$ and $x^{OE}(\underline{s}_t)$ are continuous in probabilities $\hat{\pi}(s_t | \underline{s}_{t-1})$ that satisfy Assumption 1 for all t and \underline{s}_{t-1} , where $c^{OE}(\underline{s}_t) := c^*(\underline{s}_t, \{\hat{\pi}^{OE}\})$ and $x^{OE}(\underline{s}_t) := x^*(\underline{s}_t, \{\hat{\pi}^{OE}\})$. This follows from the continuity of expected felicity in probabilities and controls, Assumption 2, and the compactness of probability spaces. For less regular problems optimal expectations may or may not exist. As to uniqueness, optimal beliefs need not be unique, as will be clear from the subsequent use of this concept.

Beliefs impact well-being directly through anticipation of future flow utility and indirectly through their effects on agent behavior. Optimal beliefs trade-off the incentive to be optimistic in order to increase expected future utility against the costs of poor outcomes that result from decisions made based on optimistic beliefs.

How does this trade-off occur in practice? Think of people as first approaching problems with optimism (“This paper will be easy to write”). But people sometimes choose not to simply accept their initial beliefs, and instead allocate cognitive resources to the problem – asking themselves whether the probabilities of a good outcome are really as high as they would like to believe (“Am I sure writing this paper will not stretch over years?”). As cognition is applied, probability assessments become more rational. We posit that the amount of cognition is directly related to the true risks and rewards of biased versus rational beliefs (“I am hesitant to commit to present the paper next week when I may not have results – let me think about it”). This description is consistent with the view that human behavior is mostly determined by the rapid and unconscious processing of the limbic system, but that for important decisions people rely more on the slower, conscious processing of the cortex. This description also matches many psychological experiments that find that agents report optimistic probabilities particularly when these probabilities or their reports do not affect payoffs. Probabilities tend to be more accurate and beliefs more rational when agents have more to lose from biased beliefs.⁵ Our optimal expectations framework is a simple model that captures this (and other) complex (and speculative) brain processes.

We view these processes – the mapping from objective to subjective probabilities – as hard-wired and subconscious, not conscious. Thus, while the interaction between optimistic and rational forces can be viewed as a model of a divided self, agents are

⁵Lichtenstein, Fischhoff, and Phillips (1982) surveys evidence on people’s overconfidence. Professionals such as weather forecasters or those who produce published gambling odds make very accurate predictions. Note also that the predictions of professionals do not seem to be due to learning from repetition (Alpert and Raiffa (1982)).

unaware of this division and of the fact that their beliefs may be biased. This lack of self-awareness implies that agents are unable to figure out the true probabilities from the model and their subjective beliefs.

So far we have focused on the optimization problem of a single agent. In a competitive economy each agent faces this maximization problem taking as given his beliefs and the stochastic process of payoff-relevant aggregate variables. In our notation, x_t^i includes the payoff-relevant variables that agent i takes as given, and so reflects the actions of all other agents in the economy. Each agent's beliefs maximize equation (4), where the states and controls are indexed by i , taking the actions of the other agents as given. In equilibrium, markets clear.

Definition 2 *A competitive optimal expectations equilibrium is a set of beliefs for each agent and an allocation such that*

- (i) each agent has optimal expectations, taking as given the stochastic process for aggregate variables;*
- (ii) each agent maximizes (1) subject to constraints taking as given his beliefs and the stochastic process for aggregate variables;*
- (iii) markets clear.*

Intuitively, optimal beliefs of each agent take as given the aggregate dynamics, and the optimal actions take as given the perceived aggregate dynamics.

2.3 Discussion

Before proceeding to the application of optimal expectations, it is worth emphasizing several points.

First, because probabilities, $\hat{\pi}^{OE}(s_t|\underline{s}_{t-1})$, are chosen once and forever, the law of iterated expectations holds with respect to the subjective probability measure and standard dynamic programming can be used to solve the agent's optimization problem. An alternative interpretation of optimal conditional probabilities is instead that the agent is endowed with optimal priors over the state space, $\hat{\pi}^{OE}(\underline{s}_T)$, and learns and updates over

time according to Bayes' rule.⁶ Thus, agents are completely “Bayesian” rational given what they know about the economic environment.

Second, optimal expectations are those that maximize well-being. The argument that is traditionally made for the assumption of rational beliefs – that such beliefs lead agents to the best outcomes – is correct only if one assumes that expected future utility flows do not affect present felicity. This is a somewhat inconsistent view: one part of the agent makes plans that trade off present and expected future utility flows, while another part of the agent actually enjoys utils but only from present consumption.⁷ Under the Jevonian view that an agent who cares about the future has felicity that depends on expectations about the future, optimal expectations give agents the highest average lifetime utility level.

To recast this point, we can ask what objective function for beliefs would make rational expectations optimal. In the general framework, this is the case if well-being counts only the felicity of the agent in the last period, so that $\mathcal{W} = E [U (c_1^*, c_2^*, \dots, c_T^*)]$. Alternatively, in the canonical time-separable model, this is the case if the objective function for beliefs omits anticipatory and memory utility, so that $\mathcal{W} = E \left[\frac{1}{T} \sum_{\tau=1}^T \beta^{\tau-1} u (c_\tau^*) \right]$.

Third, this discussion also makes clear why well-being, \mathcal{W} , uses the objective expectations operator. Optimal beliefs are not those that maximize the agent's happiness only in the states that the agent views as most likely. Instead, optimal beliefs maximize the happiness of the agent on average, across repeated realizations of uncertainty. The objective expectation captures this since the actual unfolding of uncertainty over the agent's life is determined by objective probabilities.

Fourth, the only reason for belief distortion is that current felicity depends on expected future utility flows. This is because changes in actions caused by belief distortion reduce

⁶The interpretation of the problem in terms of optimal priors requires that one specify agent beliefs following zero subjective probability events, situations in which Bayes rule provides no restrictions.

⁷See Loewenstein (1987) and the discussion of the Samuelsonian and Jevonian views of utility in Caplin and Leahy (2000).

well-being. Under rational expectations, the objective function for beliefs, \mathcal{W} , is identical to the objective function of the agent, $E[U]$. Thus, fixing beliefs to be rational, the actions of the agent maximize well-being.

To clarify this point, consider a generalized version of current felicity at time t with time-separable utility and exponential discounting

$$\hat{E}[U_t(c_{t-1}, c_t, \dots, c_T) | \mathfrak{S}_t] = \beta^{t-1} \left(\sum_{\tau=1}^{t-1} \delta^\tau u(c_{t-\tau}) + u(c_t) + \hat{E} \left[\sum_{\tau=1}^{T-t} \beta^\tau u(c_{t+\tau}) | \mathfrak{S}_t \right] \right), \quad (5)$$

where the agent discounts past utility flows at rate δ , $0 \leq \delta \leq 1/\beta$. If $\delta = 1/\beta$, then this example fits into the framework we have assumed so far; we refer to this case as *preference consistency*. If $\delta < 1$, the agent's memory utility decays through time, which has more intuitive appeal. However, in this case, an agent's ranking of utility flows across periods is not time-invariant under rational expectations (Caplin and Leahy (2000)). Thus, there is an incentive to distort beliefs in order to distort actions so as to increase well-being. In Section 5 we assume time-separable utility and exponential discounting. While the behavior of agents depends on δ , the qualitative behavior characterized by our propositions holds for any $\delta \leq 1/\beta$.

Fifth, one might be concerned that agents with optimal expectations might be driven to extinction by agents with rational beliefs. But evolutionary arguments need not favor rational expectations. Since optimal expectations respond to the costs of mistakes, agents with optimal expectations are harder to exploit than agents with fixed biases. Further, many economic environments favor agents who take on more risk (DeLong, Shleifer, Summers, and Waldmann (1990)). Finally, from a longer-term perspective and consistent with our choice of \mathcal{W} , there is a biological link between happiness and better health (Kiecolt-Glaser, McGuire, Robles, and Glaser (2002) and Cohen, Doyle, Turner, Alper, and Skoner (2003)).

Before turning to the applications, we discuss three generalizations of our approach. First, optimal expectations could be derived from a more general objective function than

a simple time-average of felicities. In particular, an earlier version of this paper defined well-being as a weighted average of the agent’s felicities.

Second, optimal subjective probabilities are chosen without any direct relation to reality. This frictionless world provides insight into the behaviors generated by the incentive to look forward with optimism when belief distortion is limited by the costs of poor outcomes. In fact, it may be that beliefs cannot be distorted far from reality for additional reasons. At some cost in terms of tractability, the frictionless model can be extended to include constraints that penalize larger distortions from reality. Beliefs would then bear some relation to reality even in circumstances in which there are no costs associated with behavior caused by distorted beliefs.

What sort of restrictions might be reasonable to impose? One could require that belief distortions be restricted to be “smooth” or lie on a coarser partition of the probability space, so that belief distortion are similar for states with similar outcomes. Alternatively, one could restrict the set of feasible beliefs to be consistent with a set of parsimonious models. For example, the agent might only be able to bias beliefs through his belief about his own ability level. Or we might require that the agent believes that his income process is some first-order Markov process rather than allow belief distortions to be completely history dependent.⁸

Finally, returning to the first point of our discussion, we maintain the assumption that conditional probabilities are fixed through time. As an alternative, one might consider beliefs as being reset in each period to maximize well-being given the new information that has arrived. We describe the relationship between these different approaches at the end of Section 5.2.

⁸If the agent were aware that his prior/model is chosen from a set of parsimonious models, then he might question these beliefs. In this case, it would make sense to impose the additional restriction that only priors for which the agent cannot detect the misspecification can be chosen, an approach being pursued in the literature on robust control. By not restricting the choice set over priors we avoid these complications.

3 Portfolio choice: optimism and a preference for skewness

In this section we consider a two-period investment problem in which an agent chooses between assets in the first period and consumes the payoff of the portfolio in the second period. We show that the agent is optimistic about the payout of his own investment and prefers positively skewed returns. The subsequent section places a continuum of these agents into a general equilibrium model with no aggregate risk, and shows that agents disagree about the returns of assets.

3.1 Portfolio choice given beliefs

There are two periods and two assets. In period one, the agent allocates his unit endowment between a risk-free asset with gross return R and a risky asset with gross return $R + Z$ (Z is the excess return of the risky asset over the risk-free rate). In period two, the agent consumes the payoff from his first-period investment.

In period one, the agent chooses his portfolio share, α , to invest in the risky asset in order to maximize felicity in the first period, $\hat{E}[U(c)]$,

$$\begin{aligned} \max_{\alpha} \quad & \beta \sum_{s=1}^S \hat{\pi}_s u(c_s), \\ \text{s.t.} \quad & c_s = R + \alpha Z_s, \\ & c_s \geq 0, \end{aligned}$$

where $u(\cdot)$ is the utility function over consumption, $u' > 0$, $u'' < 0$, $u'(0) = \infty$ and $u(0) := \lim_{c \searrow 0} u(c)$. The second constraint is set by the market. Since consumption cannot be negative, the constraint follows from the market requiring the agent to be able to meet his payment obligations in all future states.

Uncertainty is characterized by S states with ex post excess return Z_s and probabilities $\pi_s > 0$ for $s = 1, \dots, S$. Let the states be ordered so that the larger the state, the larger

the payoff, $Z_{s+1} > Z_s$, $Z_1 < 0 < Z_S$, and $Z_s \neq Z_{s'}$ for $s \neq s'$. Beliefs are given by $\{\hat{\pi}_s\}_{s=1}^S$ satisfying Assumption 1.

Noting that the second constraint can only bind for the highest or lowest payoff state, the agent's problem can be written as a Lagrangian with multipliers λ_1 and λ_S ,

$$\max_{\alpha} \beta \sum_{s=1}^S \hat{\pi}_s u(R + \alpha Z_s) - \lambda_1 (R + \alpha Z_1) - \lambda_S (R + \alpha Z_S).$$

The necessary conditions for an optimal α are

$$\begin{aligned} 0 &= \sum_{s=1}^S \hat{\pi}_s u'(R + \alpha^* Z_s) Z_s - \lambda_1 Z_1 - \lambda_S Z_S, \\ 0 &= \lambda_1 (R + \alpha^* Z_1), \\ 0 &= \lambda_S (R + \alpha^* Z_S). \end{aligned}$$

It turns out that optimal beliefs are never such that $c_s = 0$ (or $R + \alpha^* Z_s = 0$) for any s . To see this, suppose that $R + \alpha^* Z_s = 0$ for some s and consider an infinitesimal change in probabilities that results in an increase of consumption in this state. Since $u'(0) = \infty$, this causes an infinite marginal increase in well-being. Thus, optimal expectations imply $R + \alpha^* Z_s \neq 0$ for all s . By complementary slackness, $\lambda_s = 0$ for all s , and the optimal portfolio is uniquely determined by

$$0 = \sum_{s=1}^S \hat{\pi}_s u'(R + \alpha^* Z_s) Z_s \Rightarrow \alpha^* (\{\hat{\pi}\}). \quad (6)$$

3.2 Optimal beliefs

Optimal beliefs are a set of probabilities that maximize well-being, the expected time-average of felicities in the first and second period:

$$\mathcal{W} = \frac{1}{2} E \left[\hat{E}_1 [U(c^*)] + \hat{E}_2 [U(c^*)] \right].$$

In period one, the agent's felicity is the subjectively expected (anticipated) utility flow in the future period; in period two, the agent's felicity is the utility flow from actual

consumption. Substituting for our utility function and consumption and writing out the expectations, $\{\hat{\pi}^{OE}\}$ solve

$$\max_{\hat{\pi}} \frac{1}{2}\beta \sum_{s=1}^S \hat{\pi}_s u(R + \alpha^* (\{\hat{\pi}\}) Z_s) + \frac{1}{2}\beta \sum_{s=1}^S \pi_s u(R + \alpha^* (\{\hat{\pi}\}) Z_s),$$

subject to the restrictions on probabilities (Assumption 1) and where $\alpha^* (\{\hat{\pi}\})$ is given implicitly by equation (6).

To characterize optimal beliefs, first note that $\hat{\pi}_s^{OE} > 0$ for at least one state s' with $Z_{s'} < 0$ and one state s'' with $Z_{s''} > 0$. If this were not the case, the agent would view the risky asset as a money pump, and would invest or short as much of the asset as possible, so that $c_s = 0$ for $s = 1$ or for $s = S$, which contradicts our previous argument. Now consider the first-order condition associated with moving $d\hat{\pi}$ from the low-payoff state s' to the high-payoff state s'' , where both states have positive subjective probability. By the envelope condition, small changes in portfolio choice from the optimum caused by small changes in subjective probabilities lead to no change in expected utility, so that this condition is

$$\frac{1}{2}\beta (u_{s''} - u_{s'}) = -\frac{1}{2}\beta \sum_{s=1}^S \pi_s u'(R + \alpha^* Z_s) Z_s \frac{d\alpha^*}{d\hat{\pi}}, \quad (7)$$

where $u_{s'} := u(R + \alpha^* Z_{s'})$. The left-hand side is the marginal gain in ‘anticipatory utility’ in the first period from increasing $\hat{\pi}_{s''}$ at the expense of $\hat{\pi}_{s'}$; the right-hand side is the marginal loss in expected utility in the second period from the resultant change in the portfolio share of the risky asset. At the optimum, the gain in anticipatory utility balances the costs of distorting actual behavior.

Let α^{RE} denote the optimal portfolio choice for rational beliefs. The following proposition, proved in the appendix, states that the agent with optimal expectations is optimistic about the payout of his portfolio. Further, the agent with optimal expectations either takes a position opposite that of the agent with rational beliefs or is more aggressive – investing even more if the rational agent invests, or shorting more if the rational agent shorts.

Proposition 1 (*Excess risk taking due to optimism*)

(i) *Optimal belief on average are biased upwards (downwards) for states in which an agent's chosen portfolio payout is positive (negative):*

$$\text{if } \alpha^{OE} > 0, \sum_{s=1}^S (\hat{\pi}_s - \pi_s) u'(R + \alpha^{OE} Z_s) Z_s > 0;$$

$$\text{if } \alpha^{OE} < 0, \sum_{s=1}^S (\hat{\pi}_s - \pi_s) u'(R + \alpha^{OE} Z_s) Z_s < 0.$$

(ii) *An agent with optimal expectations invests more aggressively than an agent with rational expectations or in the opposite direction:*

$$\text{if } E[Z] > 0, \text{ then } \alpha^{RE} > 0, \text{ and } \alpha^{OE} > \alpha^{RE} \text{ or } \alpha^{OE} < 0;$$

$$\text{if } E[Z] < 0, \text{ then } \alpha^{RE} < 0, \text{ and } \alpha^{OE} < \alpha^{RE} \text{ or } \alpha^{OE} > 0;$$

$$\text{if } E[Z] = 0 \text{ and } S > 2, \text{ then } \alpha^{RE} = 0 \text{ and } \alpha^{OE} \neq 0.$$

To understand the first part of the proposition, note that $u'_s > 0$ for all s , and Z_s is positive for large s and negative for small s . Thus, when the agent is investing in the asset ($\alpha^{OE} > 0$), optimal expectations on average bias up the subjective probability for positive excess return states at the expense of negative excess return states.

The second part of the proposition characterizes behavior. When $E[Z] > 0$, the rational agent chooses $\alpha^{RE} > 0$ since the expected excess return on the risky asset is positive. Starting from rational beliefs, again consider a small increase in the probability of state s'' at the expense of state s' where $s'' > s'$. The left-hand side of equation (7) shows that this leads to a first-order gain in anticipatory utility. The marginal cost of this distortion, shown on the right-hand side of equation (7), is zero because the cost of a small change in portfolio allocation away from the rational optimum is only of second-order, as shown in equation (6). Thus, $\{\hat{\pi}^{OE}\} \neq \{\pi\}$ and $\alpha^{OE} \neq \alpha^{RE}$.

Further, the individual either invests more than the rational agent in the risky asset or shorts the risky asset for $E[Z] > 0$, and vice versa for $E[Z] < 0$. Why would the agent take a position in the opposite direction to the rational agent, when this implies that he is taking a negative expected payoff gamble? This occurs when anticipatory utility in the contrarian position is sufficiently large. For many utility functions, this is the case when the asset has the properties similar to a lottery ticket, that is when the asset is skewed in the opposite direction of the mean payoff.

To illustrate this point, consider a world with two states and an asset with negative

expected excess payoff, $E[Z] =: \mu_Z < 0$. We specify the payoffs Z_1 and Z_2 , such that, as we vary probabilities the mean and variance, σ_Z^2 , stay constant, but skewness decreases in π_2 .

State	Probability	Excess Payoff
1	$1 - \pi_2$	$Z_1 = \mu_Z - \sigma_Z \sqrt{\frac{\pi_2}{1 - \pi_2}}$
2	π_2	$Z_2 = \mu_Z + \sigma_Z \sqrt{\frac{1 - \pi_2}{\pi_2}}$

When π_2 is small, the asset is similar to a real-world lottery: the asset yields a small negative return with high probability and a large positive return with low probability.

Proposition 2 (*Preference for skewness*)

For unbounded utility functions, there exists a $\bar{\pi}_2$ such that for all $\bar{\pi}_2 > \pi_2 > 0$ (i) the agent is optimistic about the asset, $\hat{\pi}_2^{OE} > \pi_2$, and (ii) invests in the asset, $\alpha^{OE} > 0$, even though $E[Z] < 0$.

If the agent were to short the asset when π_2 is close to zero, $\hat{\pi}_2 < \pi_2$, and so $\pi_2 - \hat{\pi}_2$ is near zero – subjective beliefs are necessarily near rational beliefs – and $\alpha^*(\{\hat{\pi}\})$ is near α^{RE} . However, in this case, if the agent instead is optimistic about the payoff of the risky asset, $\hat{\pi}_1 < \pi_1$, then he can invest in the asset and dream about the asset paying off well. In fact, for π_1 near unity, $\hat{\pi}_1^{OE} < \pi_1$ and α^{OE} is positive. This type of behavior – buying stochastic assets with negative expected return and positive skewness – is widely observed in gambling and betting.

When $E[Z] = 0$, the cost and benefit of a marginal change in beliefs from rational beliefs are both of second order. However, the gains in anticipatory utility still dominate the costs, as long as $S > 2$. Hence, $\{\hat{\pi}^{OE}\} \neq \{\pi\}$ and the agent with optimal beliefs holds or shorts an asset that a rational agent would not.⁹ An implication is that, from the perspective of objective probabilities, the agent with optimal expectations

⁹When there are only two states and $E[Z] = 0$, we know of one special case for which $\alpha^{OE} = \alpha^{RE}$; in this case, skewness is zero (both states are equally likely). We thank Erzo Luttmer and Christian Gollier for this example.

holds an underdiversified portfolio. That is, there exists a portfolio with the same objective expected return and less objective risk since $E [R + \alpha^{OE} Z_s] = E [R + \alpha^{RE} Z_s]$ but $Var [R + \alpha^{OE} Z_s] > Var [R + \alpha^{RE} Z_s]$.

4 General equilibrium: endogenous heterogenous beliefs

In this section we place the portfolio choice problem into an exchange economy with identical agents and no aggregate risk. In an optimal expectations equilibrium, agents choose to hold idiosyncratic risk and gamble against one another, even though perfect consumption insurance is possible. These features match stylized facts about asset markets. In addition, the price of the risky asset may differ from that in an economy populated by agents with rational beliefs.

The economy consists of a continuum of agents of mass one with the same utility function and facing the same investment problem as in the previous section. As before, there are two periods, with S states in the second period. There is one technology, bonds, that is risk free and gives normalized gross return 1 (R equals unity). There is also an asset in zero net supply, equity, that gives random gross return $1 + Z$ with realized returns $1 + Z_s = \frac{1 + \varepsilon_s}{P}$ where P is the equilibrium price of equity and $\varepsilon_1 < \varepsilon_2 < \dots < \varepsilon_S$. Each agent i is initially endowed with $e = 1$ bonds. Since equity is in zero net supply, $C_s = \int_i c_s^i di = e$ in all states s .

Agent i takes her beliefs, $\{\hat{\pi}_s^i\}$, and the price of equity, P , as given, and chooses her portfolio to maximize expected utility,

$$\begin{aligned} \max_{\alpha^i} & \beta \sum_{s=1}^S \hat{\pi}_s^i u(c_s^i), \\ \text{s.t.} & c_s^i = e (1 + \alpha^i Z_s), \\ & c_s^i \geq 0. \end{aligned}$$

The first-order conditions for portfolio choice deliver a unique optimal portfolio share:

$$0 = \sum_{s=1}^S \hat{\pi}_s^i u' (1 + \alpha^* Z_s) Z_s \Rightarrow \alpha^* (\{\hat{\pi}^i\}). \quad (8)$$

Optimal beliefs maximize the well-being of each agent

$$\max_{\hat{\pi}^i} \frac{1}{2} \beta \sum_{s=1}^S \hat{\pi}_s^i u (c_s^* (\{\hat{\pi}^i\})) + \frac{1}{2} \beta \sum_{s=1}^S \pi_s u (c_s^* (\{\hat{\pi}^i\})),$$

subject to the restrictions on probabilities (Assumption 1), where $c_s^* (\{\hat{\pi}^i\}) = e (1 + \alpha^* (\{\hat{\pi}^i\}) Z_s)$ and $\alpha^* (\{\hat{\pi}^i\})$ is given by equation (8). Note that, since $Z_s = \frac{1+\varepsilon_s}{P} - 1$, optimal beliefs and asset demand depend on P .

An optimal expectations equilibrium is a set of beliefs and an allocation of assets characterized by each agent holding beliefs that maximize her well-being subject to constraints, and market clearing. Letting OE denote values in an optimal expectations equilibrium (e.g. $Z^{OE} = \frac{1+\varepsilon}{P^{OE}} - 1$) and RE denote values in a rational expectations equilibrium, we have the following proposition.

Proposition 3 (*Heterogeneous beliefs and gambling*)

- (i) an optimal expectations equilibrium exists;
- (ii) for $S > 2$, agents have heterogenous priors such that some agents hold the risky asset and some agents short the risky asset:

- there exists a subset of the agents, \mathcal{I} , such that for all $i \in \mathcal{I}$, $j \notin \mathcal{I}$, $\{\hat{\pi}^{OE,i}\} \neq \{\hat{\pi}^{OE,j}\}$ and $\hat{E}^i [Z^{OE}] > 0$, $\alpha^{OE,i} > 0$, and $\hat{E}^j [Z^{OE}] < 0$, $\alpha^{OE,j} < 0$,
- $\alpha^{OE,i} \neq \alpha^{RE} = 0$, $\{\hat{\pi}^{OE,i}\} \neq \{\pi\}$ for all i .

Since there is no aggregate risk, in the rational expectations equilibrium, no agent holds any of the risky asset and all agents have the same consumption in all states. In contrast, in an optimal expectations equilibrium agents have heterogeneous beliefs and some agents hold the asset and some short it. Consequently, agents gamble against each other and bear consumption risk.

To gain intuition for this result, consider the following example with only two states: $u(c) = \frac{1}{1-\gamma} c^{1-\gamma}$ with $\gamma = 3$, $\pi_1 = 0.25$, $\pi_2 = 0.75$, $\varepsilon_1 = -0.6$, $\varepsilon_2 = 0.2$. We choose the risky asset to have negative skewness, like returns on the US stock market. The rational

expectations equilibrium has $P^{RE} = 1$ so that $E[Z] = 0$ and no agent holds the risky asset. At this price, because the payoff of the asset is negatively skewed, agents with optimal expectations would be pessimistic about the payout of the asset and short the asset. This can be seen in Figure 1; the dotted line plots well-being as a function of $\hat{\pi}_2$ for the rational expectations price, $P = 1$. At this price, the market for the risky asset does not clear because demand is too low.

At lower prices, $E[Z] > 0$, and Proposition 1 implies that agents with optimal expectations either hold more of the asset than the agent with rational expectations or short the asset. If the price were far below P^{RE} , then the asset would have such a high expected return that it would be optimal for all agents to be optimistic about the return on the asset and to hold the asset (the dashed line in Figure 1), so again the market would not clear. The unique optimal expectations equilibrium occurs at a price of 0.986. At this equilibrium price, each agent holds one of two beliefs, each of which gives the same level of well-being. These correspond to the two local maxima of the solid line in Figure 1. One set of agents has optimistic beliefs about the return on the asset and holds the asset ($\hat{\pi}_2^{OE,i} = 0.82$ and $\alpha^{OE,i} = 0.19$); the remaining agents have pessimistic beliefs and short the asset ($\hat{\pi}_2^{OE,j} = 0.38$ and $\alpha^{OE,j} = -0.67$). The market for the risky asset clears when 78 percent of the agents are optimistic and the remaining 22 percent are pessimistic. No agents hold rational beliefs.

From an economic perspective, the most interesting result in this example is that the optimal expectations equilibrium has a 1.4 percent higher equity premium than the rational expectations equilibrium. In the example, the equity premium decreases with the skewness of the asset. For the case in which the asset is positively skewed, by symmetry of the problem, $P^{OE} > P^{RE}$. For the knife-edge case in which the asset is not skewed ($\pi_1 = 0.5$ and $\varepsilon_1 = -\varepsilon_2$), agents hold rational expectations and the optimal expectations price is equal to the rational expectations price. But this result is quite specific to this example, since by Proposition 1, this is not the case if $S > 2$ or if there is aggregate risk.




Figure 1: Well-being \mathcal{W} as a function of subjective beliefs

This relationship between skewness and expected returns is also found in the payoffs and probabilities in pari-mutuel betting at horse tracks. As in our example, in pari-mutuel betting there is no aggregate risk and there are risky assets. The longer the odds on a horse, the more positively skewed is the payoff. Golec and Tamarkin (1998) documents that the longer the odds the lower the expected return on the bet, or equivalently, the higher the price of the asset. To summarize, when returns are positively skewed, as in pari-mutuel betting, our model predicts lower expected returns than the rational model; when returns are negatively skewed, as in the US stock market, our model predicts higher expected returns than the rational model.

It is also worth noting that in an optimal expectations equilibrium there is significant trading volume, while there is no trade in the rational expectations equilibrium. Further, there is more trading when the asset is more skewed. This is consistent with the empirical findings in Chen, Hong, and Stein (2001). We note the caveat that the choice of endowments drives this result. Finally, let us speculate about a setting with heterogeneous endowments. In such a setup, it is natural to pick an equilibrium in which agents

are more optimistic about the payoffs of their initial endowments. This choice minimizes trading. Interpreting endowment as labor income, the model suggests that most agents are optimistic about the performance of the companies they work in or the countries they live in. Hence, investors overinvest in the equity of their employer and of their country relative to the predictions of standard rational models, consistent with the data (see for example Poterba (2003) on pension underdiversification and Lewis (1999) on the home bias puzzle).

5 Consumption and saving over time: undersaving and overconfidence

This section considers the behavior of an agent with optimal expectations in a multi-period consumption-saving problem with stochastic income and time-separable, quadratic utility. We show that the agent overestimates the mean of future income and underestimates the uncertainty associated with future income. That is, the agent is both unrealistically optimistic and overconfident. This is consistent with survey evidence that shows that growth rate of expected consumption is greater than that of actual consumption.

5.1 Consumption given beliefs

In each period $t = 1, \dots, T$ the agent chooses consumption to maximize the expected discounted value of utility flows from consumption subject to a budget constraint.

$$\begin{aligned} \max_{c_t} \hat{E} \left[\sum_{\tau=0}^{T-t} \beta^\tau u(c_{t+\tau}) \mid \underline{y}_t \right], \\ \text{s.t. } \sum_{\tau=0}^{T-t} R^{-\tau} (c_{t+\tau} - y_{t+\tau}) = A_t, \end{aligned}$$

where $u(c_{t+\tau}) = ac_{t+\tau} - \frac{b}{2}c_{t+\tau}^2$, initial wealth $A_1 = 0$, $a, b > 0$, $\beta R = 1$, and \underline{y}_t denotes the history of income realizations up to t . The agent's felicity at time t is given by equation (5), so that the agent has time-separable utility and discounts the future and the past

exponentially. Equation (5) allows the rate at which past utility flows are discounted, δ , to differ from the inverse of the rate at which future flows are discounted, β . We note again that the choice of δ does not affect the agent's actions given beliefs.

The only uncertainty is over income, y_t . Income has cumulative distribution function $\Pi(y_t | \underline{y}_{t-1})$ with support $[\underline{y}, \bar{y}]$ and $d\Pi(y_t) > 0$ for all $y \in Y$ where $0 < \underline{y} < \bar{y} < \frac{a}{bT}$. We assume income is independently distributed over time, and so $\Pi(y_t | \underline{y}_{t-1}) = \Pi(y_t)$. Agents however can believe that income is serially dependent, so subjective distributions are denoted by $\hat{\Pi}(y_t | \underline{y}_{t-1})$.

Assuming an interior solution, the necessary conditions for an optimum imply the Hall random walk result for consumption, but for subjective beliefs

$$c_t^* = \hat{E} \left[c_{t+1}^* | \underline{y}_t \right]. \quad (9)$$

Substituting back into the budget constraint gives the optimal consumption rule

$$c_t^* = \frac{1-R^{-1}}{1-R^{-(T-t)}} \left(A_t + y_t + \sum_{\tau=1}^{T-t} R^{-\tau} \hat{E} \left[y_{t+\tau} | \underline{y}_t \right] \right). \quad (10)$$

Optimal consumption depends on subjective expectations of future income, and on the history of income realizations through A_t . Because quadratic utility exhibits certainty equivalence in the optimal choice of consumption, the subjective variance (and higher moments) of the income process are irrelevant for the optimal consumption-saving choices of the agent, given the subjective expectation of future income.

5.2 Optimal beliefs

Optimal expectations maximize well-being subject to the agent's optimal behavior given beliefs and the restrictions on expectations. Assumption 1' in Appendix A states the restrictions on expectations for a continuous state space. We incorporate optimal behavior directly into the objective function and characterize consumption choices, $\{c_t^{OE}\}$, implied by optimal beliefs. Optimal beliefs, \hat{E}^{OE} and $\{\hat{\Pi}^{OE}\}$, implement these consumption

choices given optimal behavior on the part of the agent.¹⁰

Since the objective is a sum of utility functions, it is concave in future consumption. And since the agent's behavior depends only on the subjective certainty-equivalent of future income, optimal beliefs minimize subjective uncertainty. Thus, future income is optimally perceived as certain, which is an extreme form of overconfidence.

Using the fact that optimal beliefs are certain and the consumption Euler equation, $\hat{E} \left[u(c_{t+\tau}^*) | \underline{y}_t \right] = u \left(\hat{E} \left[c_{t+\tau}^* | \underline{y}_t \right] \right) = u(c_t^*)$, so that the agent's felicity at time t can be written as

$$\hat{E} \left[U_t(c_1^*, c_2^*, \dots, c_T^*) | \underline{y}_t \right] = \beta^{t-1} \left(\sum_{\tau=1}^{t-1} \delta^\tau u(c_{t-\tau}^*) + u(c_t^*) \sum_{\tau=t}^T \beta^{\tau-t} \right).$$

Subjective expectations are chosen to yield the path of $\{c_t^*\}$ that maximizes well-being

$$\frac{1}{T} E \left[\underbrace{u(c_1^*) \sum_{\tau=1}^T \beta^{\tau-1}}_{\hat{E}[U_1^* | \underline{y}_1]} + \underbrace{\beta \left(\delta u(c_1^*) + u(c_2^*) \sum_{\tau=2}^T \beta^{\tau-2} \right)}_{\hat{E}[U_2^* | \underline{y}_2]} + \dots + \underbrace{\beta^{T-1} \left(\sum_{\tau=1}^{T-1} \delta^\tau u(c_{T-\tau}^*) + u(c_T^*) \right)}_{\hat{E}[U_T^* | \underline{y}_T]} \right],$$

subject to the budget constraint. Collecting terms, the objective simplifies to

$$\frac{1}{T} E \left[\sum_{t=1}^T \psi_t u(c_t^*) \right], \quad (11)$$

where $\psi_t = \beta^{t-1} \left(1 + \sum_{\tau=1}^{T-t} (\beta^\tau + (\beta\delta)^\tau) \right)$. Notice that, regardless of δ , the average consumption path of agents is not constant. Only if the objective for beliefs were to ignore anticipatory utility and memory utility ($\delta = 0$) so that $\psi_t = \beta^{t-1}$, would beliefs be rational and the expected consumption path standard.

¹⁰In taking this approach, we are assuming that the optimal choice of consumption and thus $\hat{E}^{OE} \left[y_{t+\tau} | \underline{y}_t \right]$ does not require violation of the assumptions on expectations, which can be checked. That is, if the support of y_t is small, belief distortion may be constrained by the range of possible income realizations. To incorporate these constraints directly, one would solve for optimal $\hat{E} \left[y_{t+\tau} | \underline{y}_t \right]$ by replacing $c^* \left(\underline{y}_t, \left\{ \hat{\Pi}^{OE} \right\} \right)$ using equation (10) and impose Assumption 1'.

Under optimal expectations, the first-order condition implies that expected consumption growth between t and $t + \tau$ is given by

$$u' (c_t^{OE}) = \frac{\psi_{t+\tau}}{\psi_t} R^\tau E \left[u' (c_{t+\tau}^{OE}) | \underline{y}_t \right],$$

which, substituting for the quadratic utility function, implies that

$$c^{OE} (\underline{y}_t) = \frac{a}{b} - \frac{\psi_{t+\tau}}{\psi_t} R^\tau \left(\frac{a}{b} - E \left[c^{OE} (\underline{y}_{t+\tau}) | \underline{y}_t \right] \right). \quad (12)$$

Level consumption is recovered by substituting into the budget constraint and taking objective conditional expectations.

Given this characterization of optimal behavior, agents are optimistic at every time and state. Define human wealth at t as the present value of current and future labor income, $H_t = \sum_{\tau=0}^{T-t} R^{-\tau} y_{t+\tau}$.

Proposition 4 (*Overconsumption due to optimism*)

For all $t \in \{1, \dots, T - 1\}$:

(i) on average, agents revise down their expectation of human wealth over time:

$$\hat{E}^{OE} \left[H_{t+1} | \underline{y}_t \right] > E \left[\hat{E}^{OE} \left[H_{t+1} | \underline{y}_{t+1} \right] | \underline{y}_t \right];$$

(ii) on average, consumption falls over time: $c_t^{OE} > E \left[c_{t+1}^{OE} | \underline{y}_t \right]$;

(iii) agents are optimistic about their future consumption: $\hat{E}^{OE} \left[c_{t+1}^{OE} | \underline{y}_t \right] > E \left[c_{t+1}^{OE} | \underline{y}_t \right]$.

The first point of the proposition states that agents overestimate their present discounted value of labor income and on average revise their beliefs downward between t and $t + 1$. The second point states that consumption on average falls between t and $t + 1$. Because on average the agent revises down expected future income, on average consumption falls over time. The proof follows directly from the expected change in consumption given by equation (12) and noting that $\frac{a}{b} - c_t^{OE} (\underline{y}_t) > 0$ and $\frac{\psi_{t+1}}{\psi_t} R < 1$. Finally, the optimal subjective expectation of future consumption exceeds the rational expectation of future consumption. This is optimism. Part (iii) follows from part (ii) and equation (9). In sum, households are unrealistically optimistic, and, in each period, are on average

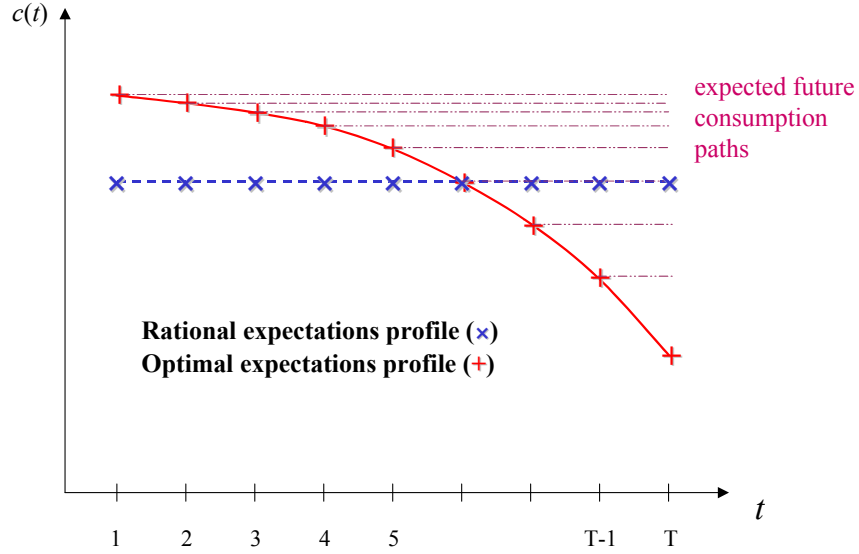


Figure 2: Average life-cycle consumption profiles

surprised that their incomes are lower than they expected, and so, on average, household consumption declines over time.

Figure 2 summarizes these results. The agent starts life optimistic about future income. At each point in time the agent expects that on average consumption will remain at the same level. Over time, the agent observes on average that income is less than he expected, and consumption typically declines over his life. Note that the agent updates his beliefs according to $\{\hat{\Pi}\}$, but does not learn over time that $\{\hat{\Pi}\}$ is incorrect because he does not know that his income is identically and independently distributed over time. The agent merely observes one realization of income at each age and (on average) believes that he was unlucky.

While quadratic utility makes this example quite tractable, the agent's overconfidence is extreme. Before each period the agent is certain about what his future income will be, and this belief is contradicted by the realization. But, as seen in equation (10), less extreme overconfidence would not alter consumption choices, only reduce the agent's

felicity early in life.

This optimism matches survey evidence on desired and actual life-cycle consumption profiles. Barsky, Juster, Kimball, and Shapiro (1997) finds that households would choose upward sloping consumption profiles. But survey data on actual consumption reveal that households have downward sloping or flat consumption profiles (Gourinchas and Parker (2002), Attanasio (1999)). In our model, households expect and plan to have constant marginal utility since $\beta R - 1 = 0$. However, on average marginal utility rises at the age-specific rate $\frac{\psi_t}{\psi_{t+1}} - 1 > 0$. Thus, in the model, the desired rate of increase of consumption exceeds the average rate of increase, as in the real world. In addition, the model matches observed household consumption behavior in that average life-cycle consumption profiles are concave – consumption falls faster (or rises more slowly) later in life.

In general, in consumption-saving problems, the relative curvatures of utility and marginal utility determine what beliefs are optimal. Uncertainty about the future enters the objective for beliefs both through the expected future *level* of utility and through the agent's behavior, which depends on expected future *marginal* utility. For utility functions with decreasing absolute risk aversion, greater subjective uncertainty leads to greater precautionary saving through the curvature in marginal utility. This has some benefit in terms of less distortion of consumption. In such cases optimal beliefs may consist of a large positive bias for both expected income and its variance.

We conclude this section by using our consumption-saving problem to make four points about the dynamic choices of agents with optimal expectations.

First, given that in expectation the consumption of the agent is always declining, the costs of optimism early in life could be extreme for long-lived agents. But, illustrating a general point, optimal expectations depend on the horizon in a way that mitigates these possible costs. The behavior of an agent with a long horizon is close to that of an agent with rational expectations. For T large but finite, an agent with optimal expectations consumes a small amount more for most his life, leading to a significant decline in consump-

tion at the end of life. As the horizon becomes infinite, at any fixed age, the consumption choice of the agent with optimal beliefs converges to that of the agent with rational beliefs as the subjective expectation of human wealth converges to the rational expectation. Formally, for any t , as $T \rightarrow \infty$, $c_t^{OE}(\underline{y}_t) \rightarrow c_t^{RE}(\underline{y}_t)$, $\hat{E}^{OE}[H_{t+1}|\underline{y}_t] \rightarrow E[H_{t+1}|\underline{y}_t]$, and $c_t^{OE}(\underline{y}_t) \rightarrow E[c_{t+1}^{OE}(\underline{y}_{t+1})|\underline{y}_t]$.¹¹ Beliefs become more rational as the stakes become larger.

Second, an agent with optimal expectations may choose not to insure future income when offered an objectively fair insurance contract. Formally, let the agent face an additional binary decision in period one: whether or not to exchange all current and future income for $B = E[H_1|\underline{y}_1]$. A rational agent would always take this contract, while the agent with optimal expectations may choose not to insure consumption. Interestingly, since beliefs affect whether the agent insures or not, the addition of the possibility of insurance may change what beliefs are optimal.

Optimal expectations are either the beliefs that maximize well-being conditional on inducing the agent to reject the insurance, or the beliefs that maximize well-being conditional on inducing the agent to accept the insurance. The former are the optimal expectations from Proposition 4. These beliefs are optimal for the problem without the constraint, and the agent rejects the insurance because both income streams are perceived as certain and $\hat{E}[H_1|\underline{y}_1] > E[H_1|\underline{y}_1] = B$. Well-being in this case is

$$\frac{1}{T} \sum_{t=1}^T \psi_t E[u(c_t^*)|\underline{y}_1]$$

from equation (11). The latter, optimal beliefs conditional on accepting the insurance, are irrelevant for well-being, provided that the agent believes that $\hat{E}[H_1|\underline{y}_1]$ is small enough

¹¹It can be seen that $c_1^{OE} \rightarrow c_1^{RE}$ by taking conditional expectations of the budget constraint and repeatedly substituting $E[c_{t+\tau}^{OE}|\underline{y}_t]$, for all τ , from the Euler equation (12) to solve for c_1^{OE} , and noting that $\frac{\psi_{t+\tau}}{\psi_t} R^\tau \rightarrow 1$ as $T \rightarrow \infty$ and $\frac{\psi_t}{\psi_{t+\tau}} R^{-\tau} < \psi_t$ for $\tau, T \rightarrow \infty$. This together with equation (10) imply that $\hat{E}^{OE}[H_2|\underline{y}_1] \rightarrow E[H_2|\underline{y}_1]$. Again using the fact that $\frac{\psi_{t+\tau}}{\psi_t} R^\tau \rightarrow 1$ as $T \rightarrow \infty$, equation (12) implies that for finite t , $c_t^{OE} \rightarrow c_t^{RE}$, so that the first two results also hold for any finite t (not just $t = 1$).

and/or the process for $\{y\}$ uncertain enough that he accepts the insurance.¹² Well-being in this case is

$$\frac{1}{T} \sum_{t=1}^T \psi_t u \left(c^{FI} \left(\underline{y}_1 \right) \right),$$

where $c^{FI} \left(\underline{y}_1 \right) = \frac{R-1}{R-R^{1-T}} E \left[\sum_{t=1}^T R^{1-t} y_t | \underline{y}_1 \right]$.

Risk determines which expectations are optimal. Well-being decreases in objective income risk when the agent rejects the insurance, while it is invariant to risk if he accepts the insurance. If objective income risk is small, then the cost of distorted beliefs – variable future consumption – is small, and optimal expectations are optimistic. If objective income risk is large, optimal expectations are more rational and induce the agent to insure his future income.

Third, at the start of life the agent facing the problem with the option to insure income may have a lower level of felicity than the agent facing the problem without this option. Informally, we might think of an agent approaching their life blithely optimistic about their future. Given no choice of insurance, this is indeed optimal. However, if placed in an environment with large amounts of income risk, an agent given the opportunity to insure considers his life more realistically, puts more weight on possible bad states of the world, and chooses insurance. Since $c^{OE} \left(\underline{y}_1 \right) > c^{FI} \left(\underline{y}_1 \right)$, the agent who has and accepts the option to insure is less happy initially.¹³

Finally, what if beliefs were chosen in each period to maximize well-being? Suppose that the agent in each period chooses his actions taking as given his own beliefs in the future, which are possibly different. The agent's felicity is the present discounted value of utility flows evaluated using his own subjective beliefs. This can be viewed as if the agent in each period is a different self that knows the conditional be-

¹²While nothing formally requires this, it seems natural to assume that expectations are rational in this case.

¹³On average, the agent who has and accepts the option to insure has greater levels of felicity later in life. This is because lifetime well-being with the option to insure is greater than or equal to lifetime well-being without the option.

liefs of his future selves. The well-being function for optimal beliefs at time t would then be $\mathcal{W}^t := E \left[\frac{1}{T} \sum_{\tau=1}^T \hat{E}^\tau [U(c_1^*, c_2^*, \dots, c_T^*) | \underline{s}_\tau] | \underline{s}_t \right]$ where $\hat{E}^\tau [\cdot | \underline{s}_\tau]$ denotes the beliefs of the agent at time τ under this alternative assumption. $\hat{E}^t [\cdot | \underline{s}_\tau]$ maximizes \mathcal{W}^t given $\left\{ \hat{E}^\tau [\cdot | \underline{s}_\tau] \right\}_{\tau > t}$ and the future decision rules that these beliefs induce.

Because the objective function changes through time, typically it is not the case that the agent updates probabilities according to Bayes' law. That is, $\hat{E}^t [\cdot | \underline{s}_\tau]$ varies across an agent's selves in different periods $t < \tau$. However, in the application of this section, an agent's selves agree.

Proposition 5 (*Time consistency of beliefs*)

In this consumption-saving problem, optimal expectations are time consistent: $\hat{E}^{OE,t} [\cdot | \underline{s}_\tau]$ is independent of t for all possible histories and $\tau \geq t$.

This result obtains here because of the extremity of overconfidence. Consider first the choice of beliefs at time t following an event at $t + s$, viewed subjectively as having zero probability. These beliefs do not influence either the actions or anticipatory utility of the agent at time t ; they only influence the actions and anticipatory utility of the agent at time $t + s$. Thus, beliefs following realizations of income besides the expected one are chosen simply to maximize the expected utility of that agent in that period. The perspective from which one chooses these optimal beliefs is irrelevant. If the income realization matches the expected level, then consumption remains constant, and the agent continues to hold the certain beliefs that they held in the previous period (by Bayes' law). Note that this argument pins down the profile of optimal income expectations, $\hat{E} [y_t]$, which increases in a pattern opposite the average consumption profile, $E [c_t]$.

6 Conclusion

This paper introduces a model of utility-serving biases in beliefs. While our applications highlight many of the implications of our theory, many remain to be explored.

First, the specification of possible events seems to be more important in a model with optimal expectations than it is in a model with rational expectations. For example, an

optimal expectations equilibrium in a world with only certain outcomes is different from the equilibrium in the same world with an available sunspot or public randomization device. With the randomization device agents can gamble against one another.

Second, agents with optimal expectations can be optimistic about uncertain events, and therefore can be better off with the later resolution of uncertainty. For instance, you tell someone that they are going to receive gifts on their birthday but you do not tell them what those gifts are until their birthday.¹⁴ More generally, because more information can change the ability to distort beliefs, agents can be better off not receiving information despite the benefits of better decision making. However, without relaxing the assumptions of expected utility theory and Bayesian updating, agents would not choose that uncertainty be resolved later because agents take their beliefs as given.

Third, we conjecture that the agent who faces the same problem again and again, and so faces the possibility of large losses from an incorrect specification of probabilities, will, in our framework, have a better assessment of probabilities. Thus, optimal expectations agents are not easy to turn into “money pumps,” although they may exhibit behavior far from that generated by rational expectations in one-shot games.

Fourth, and closely related, to what extent do optimal beliefs give an evolutionary advantage or disadvantage relative to rational beliefs? On the one hand, agents with optimal expectations make poorer decisions. On the other hand, agents with optimal expectations may take on more risk, which can lead to an evolutionary advantage.

Finally, optimal expectations has promising applications in strategic environments. In a strategic setting, each agent’s beliefs are set taking as given the reaction functions of other agents.

¹⁴A surprise party for an agent raises the possibility in the agent’s mind that he might get more surprise parties in the future and he enjoys looking forward to this possibility.

Appendixes

A Optimal expectations when the state space is continuous

In the main text, we define optimal expectations when the state space is finite and discrete. To consider random variables with continuous distributions, we extend our definitions. Let $\{\mathbb{S}_T, \mathcal{F}, \Pi\}$ denote the state space, σ -algebra, and objective probability measure. Let $\mathbb{F} = \{\mathcal{F}_0, \dots, \mathcal{F}_T\}$ be a filtration. Let $\{\hat{\Pi}\}$ and \hat{E} denote the subjective probability measure and expectation respectively. First, agent optimization given continuously distributed random variables is standard. Second, it is mathematically simpler to state the restrictions on subjective beliefs in terms of subjective conditional expectations. Thus, one solves for optimal expectations by choosing $\hat{E}[A|\mathcal{F}_t]$ for any \mathcal{F}_t in the filtration \mathbb{F} and any event $A \subseteq \mathbb{S}_T$ to maximize the functional objective and Assumption 1 is replaced by

Assumption 1' (Restrictions on probabilities for a continuous state space) For every $\mathcal{F}_t \in \mathbb{F}$

- (i) $\hat{E}[\mathbb{S}_T|\mathcal{F}_t] = 1$
- (ii) $\hat{E}[f|\mathcal{F}_t] \geq 0$ for any nonnegative function $f: \mathbb{S}_T \mapsto \mathbb{R}$ which is \mathcal{F}_t -measurable
- (iii) $\hat{E}[A|\mathcal{F}_t] = \hat{E}\left[\hat{E}[A|\mathcal{F}_{t+\tau}]|\mathcal{F}_t\right]$ for any $\tau \geq 0$ and any event A
- (iv) Π is a dominating measure of $\hat{\Pi}$.

B Proofs of Propositions

B.1 Proof of Proposition 1

(i) We prove the case for $\alpha^{OE} > 0$; the case for $\alpha^{OE} < 0$ is analogous. For $\alpha^{OE} > \alpha^{RE}$, when the asset pays off poorly, marginal utility is higher (lower) for the agent with the higher (lower) share invested in the risky asset:

$$\begin{aligned} u'(R + \alpha^{OE} Z_s) &\geq u'(R + \alpha^{RE} Z_s) \text{ for } s \text{ such that } Z_s \leq 0 \\ u'(R + \alpha^{OE} Z_s) &< u'(R + \alpha^{RE} Z_s) \text{ for } s \text{ such that } Z_s > 0 \end{aligned} \tag{B.1}$$

Combining this with the first-order condition of the agent with rational expectations,

$$\sum_{s \ni Z_s \leq 0} \pi_s u' (R + \alpha^{RE} Z_s) Z_s + \sum_{s \ni Z_s > 0} \pi_s u' (R + \alpha^{RE} Z_s) Z_s = 0,$$

yields

$$\sum_{s \ni Z_s \leq 0} \pi_s u' (R + \alpha^{OE} Z_s) Z_s + \sum_{s \ni Z_s > 0} \pi_s u' (R + \alpha^{OE} Z_s) Z_s < 0.$$

Subtracting this from the first-order condition of the agent with optimal expectations gives the desired inequality

$$\sum_{s=1}^S (\hat{\pi}_s^{OE} - \pi_s) u' (R + \alpha^{OE} Z_s) Z_s > 0 \quad (\text{B.2})$$

Thus, if we can show that $\alpha^{OE} > 0$ implies $\alpha^{OE} > \alpha^{RE}$ the proof of (i) is complete. This follows from the second point of the proposition, which we now prove.

(ii) The proof of the sign of α^{RE} in each case is standard and omitted. We first treat the case of $E[Z] > 0$ and $\alpha^{RE} > 0$, the case of $E[Z] < 0$ and $\alpha^{RE} < 0$ is analogous, and we treat $E[Z] = 0$ and $\alpha^{RE} = 0$ subsequently.

We first show that an agent with arbitrary beliefs invests more in the risky asset (or shorts it less) as the subjective probability of a state s'' with $Z_{s''} > 0$ is increased relative to a state s' with $Z_{s'} < 0$. Examine the agent's first-order condition for α^* and consider moving $d\hat{\pi}$ from s' to s''

$$0 = (u'(R + \alpha^* Z_{s''}) Z_{s''} - u'(R + \alpha^* Z_{s'}) Z_{s'}) d\hat{\pi} + \sum_{s=1}^S \hat{\pi}_s u''(R + \alpha^* Z_s) Z_s^2 d\alpha^*$$

$$\frac{d\alpha^*}{d\hat{\pi}} = -\frac{u'(R + \alpha^* Z_{s''}) Z_{s''} - u'(R + \alpha^* Z_{s'}) Z_{s'}}{\sum_{s=1}^S \hat{\pi}_s u''(R + \alpha^* Z_s) Z_s^2} > 0$$

since the denominator is negative and $Z_{s''} > 0 > Z_{s'}$.

Suppose for purposes of contradiction, that $0 < \alpha^{OE} \leq \alpha^{RE}$. As in the proof of part (i), we have

$$u'(R + \alpha^{OE} Z_s) \leq u'(R + \alpha^{RE} Z_s) \text{ for } s \ni Z_s \leq 0$$

$$u'(R + \alpha^{OE} Z_s) > u'(R + \alpha^{RE} Z_s) \text{ for } s \ni Z_s > 0$$

which implies from the first order condition of the agent with rational expectations,

$$\sum_{s=1}^S \pi_s u' (R + \alpha^{OE} Z_s) Z_s > 0. \quad (\text{B.3})$$

Now to establish the contradiction, the first-order condition for beliefs, equation (7), implies

$$\text{sign} [\beta (u_{s''} - u_{s'})] = \text{sign} \left[- \sum_{s=1}^S \pi_s u' (R + \alpha^{OE} Z_s) Z_s \frac{d\alpha^*}{d\hat{\pi}} \right]$$

From $\frac{d\alpha^*}{d\hat{\pi}} > 0$ and equation (B.3), the sign of the right hand side is strictly negative, while the fact that we are assuming $\alpha^{OE} > 0$ implies that $u_{s''} > u_{s'}$ and the left hand side is strictly positive, a contradiction. Therefore either $\alpha^{OE} > \alpha^{RE}$ or $\alpha^{OE} \leq 0$. The final step is to rule out $\alpha^{OE} = 0$. If $\alpha^{OE} = 0$ then $\beta (u_{s''} - u_{s'}) = 0$, so that by the first-order condition for beliefs

$$0 = \sum_{s=1}^S \pi_s u' (R + \alpha^{OE} Z_s) Z_s. \quad (\text{B.4})$$

But $\alpha^{OE} = 0$ cannot solve equation (B.4) because equation (B.4) is the same as the first-order condition for the optimal portfolio choice of the rational agent, and the objective of the rational agent is globally concave with a unique α^{RE} satisfying equation (B.4).

Finally, we prove that when $E[Z] = 0$ and $\alpha^{RE} = 0$, $\alpha^{OE} \neq 0$. Suppose that instead $\{\hat{\pi}^{OE}\}$ were such that $\alpha^{OE} = 0$, which occurs if and only if $\hat{E}^{OE}[Z] = 0$. These beliefs actually satisfy the first-order condition for optimal expectations because 1) there is no gain to the marginal belief distortion since $u_{s''} = u_{s'}$ and 2) starting from α^{RE} the first-order cost of a small change in optimal portfolio choice is zero. We show however, that the second order condition is violated for some beliefs such that $\hat{E}[Z] = 0$, which means that there is a deviation from this set of beliefs that increases well-being and therefore $\alpha^{OE} \neq 0$.

The second order condition for the same $d\hat{\pi}$ that moves an infinitesimal probability from s' to s'' , where $Z_{s''} > 0$, $Z_{s'} < 0$, and $\hat{\pi}_{s'} > 0$ is

$$\frac{d^2\mathcal{W}}{d\hat{\pi}d\hat{\pi}} = 2 \frac{\partial^2\mathcal{W}}{\partial\hat{\pi}\partial\alpha} \frac{d\alpha^*}{d\hat{\pi}} + \frac{\partial^2\mathcal{W}}{\partial\hat{\pi}\partial\hat{\pi}} + \frac{\partial^2\mathcal{W}}{\partial\alpha^2} \left(\frac{d\alpha^*}{d\hat{\pi}} \right)^2 + \frac{\partial\mathcal{W}}{\partial\alpha} \frac{d^2\alpha^*}{d\hat{\pi}^2}$$

Since W is linear in probabilities, $\frac{d^2\mathcal{W}}{d\hat{\pi}d\hat{\pi}} = 0$. Now, omitting $\frac{1}{2}\beta$ from all terms,

$$\begin{aligned}\frac{\partial\mathcal{W}}{\partial\alpha} &= \sum_{s=1}^S (\pi_s + \hat{\pi}_s) u'_s Z_s \\ \frac{\partial^2\mathcal{W}}{\partial\alpha^2} &= \sum_{s=1}^S (\pi_s + \hat{\pi}_s) u''_s Z_s^2 \\ \frac{\partial^2\mathcal{W}}{\partial\hat{\pi}\partial\alpha} &= u'_{s''} Z_{s''} - u'_{s'} Z_{s'}\end{aligned}$$

We evaluate this second order condition for $\{\hat{\pi}\}$ such that $\hat{E}[Z] = 0$ so that $\alpha^* = 0$ and u' and u'' are independent of s , yielding

$$\begin{aligned}\frac{d^2\mathcal{W}}{d\hat{\pi}d\hat{\pi}} &= \left\{ 2u' [Z_{s''} - Z_{s'}] + u'' \left[\sum_{s=1}^S (\pi_s + \hat{\pi}_s) Z_s^2 \right] \frac{u' [Z_{s'} - Z_{s'']}}{u'' \sum_{s=1}^S \hat{\pi}_s Z_s^2} \right\} \frac{d\alpha^*}{d\hat{\pi}} + u' \sum_{s=1}^S (\pi_s + \hat{\pi}_s) Z_s \frac{d^2\alpha^*}{d\hat{\pi}^2} \\ &= \left\{ 2 - \frac{\sum_{s=1}^S (\pi_s + \hat{\pi}_s) Z_s^2}{\sum_{s=1}^S \hat{\pi}_s Z_s^2} \right\} [Z_{s''} - Z_{s'}] u' \frac{d\alpha^*}{d\hat{\pi}} + u' \left\{ \sum_{s=1}^S \pi_s Z_s + \sum_{s=1}^S \hat{\pi}_s Z_s \right\} \frac{d^2\alpha^*}{d\hat{\pi}^2} \\ &= \left\{ 1 - \frac{\sum_{s=1}^S \pi_s Z_s^2}{\sum_{s=1}^S \hat{\pi}_s Z_s^2} \right\} [Z_{s''} - Z_{s'}] u' \frac{d\alpha^*}{d\hat{\pi}}\end{aligned}$$

where the third equality makes use of $\sum \pi_s Z_s = 0$ and $\sum \hat{\pi}_s Z_s = 0$. Thus, any $\{\hat{\pi}\}$ such that $\sum_{s=1}^S \pi_s Z_s^2 < \sum_{s=1}^S \hat{\pi}_s Z_s^2$ and $\hat{E}[Z] = 0$ has $\frac{d^2\mathcal{W}}{d\hat{\pi}d\hat{\pi}} > 0$, and so there exists a deviation that increases well-being, completing the proof. This final step requires $S > 2$; for $S = 2$, the second order condition is necessarily zero and there are cases where $\alpha^{OE} = \alpha^{RE}$. We conjecture that $\alpha^{OE} = \alpha^{RE}$ only occurs for $S = 2$, $Z_1 = -Z_2$, and $\pi_1 = \pi_2 = 1/2$.

B.2 Proof of Proposition 2

To avoid arbitrage we consider only π_1 large enough such that $Z_2 > 0$. We show that as $\pi_1 \rightarrow 1$, well-being when investing in the asset is higher than when shorting the asset. We do this by constructing a lower bound for well-being when investing in the asset ($\underline{\mathcal{W}}^+(\pi_1)$) and an upper bound when shorting the asset ($\overline{\mathcal{W}}^-(\pi_1)$) and showing that $\lim_{\pi_1 \rightarrow \infty} \underline{\mathcal{W}}^+(\pi_1) > \lim_{\pi_1 \rightarrow \infty} \overline{\mathcal{W}}^-(\pi_1)$. Define well-being as a function of subjective and objective beliefs, given optimal agent behavior, as

$$\mathcal{W}(\hat{\pi}_1; \pi_1) := \frac{1}{2}\beta(\pi_1 + \hat{\pi}_1) u(R + \alpha^* Z_1) + \frac{1}{2}\beta(2 - \pi_1 - \hat{\pi}_1) u(R + \alpha^* Z_2)$$

where $Z_s = Z_s(\pi_1)$, and $\alpha^* = \alpha^*(\hat{\pi}_1; Z_1(\pi_1), Z_2(\pi_1))$.

Step 1: $\lim_{\pi_1 \rightarrow 1} \mathcal{W}(\cdot) = \infty$ for $\alpha > 0$.

Consider an optimistic belief, $\hat{\pi}'_1$, $0 < \hat{\pi}'_1 < \pi_1$, such that the agent invests in the asset, $\alpha' > 0$. Since $\hat{\pi}'_1$ may be suboptimal, well-being with this belief is a lower bound for the well-being of the agent conditional on $\alpha > 0$. Define $\underline{\mathcal{W}}^+(\pi_1) := \mathcal{W}(\hat{\pi}'_1; \pi_1) \leq \mathcal{W}(\hat{\pi}_1^{OE}(\pi_1), \pi_1)$.

Taking the limit as skewness goes to infinity

$$\begin{aligned} \lim_{\pi_1 \rightarrow 1} \underline{\mathcal{W}}^+(\pi_1) &= \frac{1}{2}\beta(1 + \hat{\pi}'_1) \lim_{\pi_1 \rightarrow 1} u(R + \alpha'Z_1) + \frac{1}{2}\beta(1 - \hat{\pi}'_1) \lim_{\pi_1 \rightarrow 1} u(R + \alpha'Z_2) \\ &= \frac{1}{2}\beta(1 + \hat{\pi}'_1) u(R + \alpha'\mu_Z) + \frac{1}{2}\beta(1 - \hat{\pi}'_1) \lim_{\pi_1 \rightarrow 1} u(R + \alpha'Z_2) \\ &= \infty \end{aligned}$$

since $Z_1 \rightarrow \mu_Z$, $Z_2 \rightarrow \infty$ and $\lim_{c \rightarrow \infty} u(c) = \infty$.

Step 2: $\lim_{\pi_1 \rightarrow 1} \mathcal{W}(\cdot) < \infty$ for $\alpha < 0$.

Define an upper bound for well-being by choosing the portfolio and subjective beliefs subject only to the conditions that the agent shorts the asset, that the agent is pessimistic about the payout, and that the portfolio is feasible:

$$\begin{aligned} \overline{\mathcal{W}}^-(\pi_1) &: = \frac{1}{2}\beta \sup_{\alpha, \hat{\pi}_1} [(\pi_1 + \hat{\pi}_1) u(R + \alpha Z_1) + (2 - \pi_1 - \hat{\pi}_1) u(R + \alpha Z_2)] \\ \text{s.t.} \quad &\alpha < 0 \\ &\hat{\pi}_1 > \pi_1 \\ &R + \alpha Z_2 \geq 0 \end{aligned}$$

$\overline{\mathcal{W}}^-(\pi_1)$ is an upper bound since we do not restrict α to be the optimal agent's choice given $\hat{\pi}_1$. The optimal $\hat{\pi}_1 = 1$, and the first and third constraints become $-\frac{R}{Z_2} \leq \alpha \leq 0$ (which is not the null set since $Z_2 > 0$), so that this can be re-written as

$$\begin{aligned} \lim_{\pi_1 \rightarrow 1} \overline{\mathcal{W}}^-(\pi_1) &= \frac{1}{2}\beta \lim_{\pi_1 \rightarrow 1} \sup_{-\frac{R}{Z_2} \leq \alpha \leq 0} [(1 + \pi_1) u(R + \alpha Z_1) + (1 - \pi_1) u(R + \alpha Z_2)] \\ &< \frac{1}{2}\beta \lim_{\pi_1 \rightarrow 1} \left[(1 + \pi_1) u\left(R + \left(-\frac{R}{Z_2}\right) Z_1\right) + (1 - \pi_1) u(R) \right] = \beta u(R) \end{aligned}$$

where the second line follows from substituting the best portfolio choice in each state separately and $\lim_{\pi_1 \rightarrow 1} \frac{Z_1}{Z_2} = 0$.

The proof follows from $\lim_{\pi_1 \rightarrow 1} \underline{\mathcal{W}}^+(\pi_1) > \lim_{\pi_1 \rightarrow 1} \overline{\mathcal{W}}^-(\pi_1)$ and Proposition 1.

B.3 Proof of Proposition 3

(i) Write the well-being of an agent as a function of subjective beliefs and the price given $\{\pi\}$ and optimal agent behavior:

$$\mathcal{W}(\{\hat{\pi}\}, P) = \frac{1}{2}\beta \sum_{s=1}^S \hat{\pi}_s u(1 + \alpha^* Z_s) + \frac{1}{2}\beta \sum_{s=1}^S \pi_s u(1 + \alpha^* Z_s)$$

This function is well-defined for the set of prices and beliefs such that the agent chooses positive consumption in every state. Denote this set M . In M the function \mathcal{W} is continuous in prices and subjective probabilities because α^* is continuous in subjective probabilities and prices. M is not closed, but we now show that P^{OE} and all $\{\hat{\pi}^{OE}\}$ do not lie outside M or in the set $\{\text{Closure}(\mathcal{M}) \setminus \mathcal{M}\}$. First, consider prices such that $Z_1 \geq 0$ or $Z_S \leq 0$. In this case, all agents would have an identical arbitrage opportunity for any possible subjective beliefs, except possibly for $\hat{\pi}_1 = 1$ or $\hat{\pi}_S = 1$. Hence, this cannot constitute an equilibrium because agents would all choose to buy or all choose to short the risky asset, and so the market for the risky asset would not clear. Thus, the equilibrium price must lie on the interior of the set $P \in (1 + \varepsilon_1, 1 + \varepsilon_S)$. Second, consider beliefs such that an agent chooses $c_s = 0$ for some s . Because $u'(0) = \infty$ in some state, a marginal increase in $\hat{\pi}_1$ or $\hat{\pi}_S$ leads to $c_s > 0$ for all s , and results in an infinite increase in well-being.

We will argue that at a low enough price, $\mathcal{W}(\{\hat{\pi}\}, P)$ is maximized by beliefs such that $\hat{E}[Z] > 0$ and $\alpha^* > 0$, and at a high enough price, $\mathcal{W}(\{\hat{\pi}\}, P)$ is maximized by beliefs such that $\hat{E}[Z] < 0$ and $\alpha^* < 0$. Then, by continuity of $\mathcal{W}(\{\hat{\pi}\}, P)$, either at some intermediate price there are multiple global maxima, some with $\hat{E}[Z] > 0$ and some with $\hat{E}[Z] < 0$, or at some intermediate price there is a unique global maximum with $\hat{E}[Z] = 0$ and $\alpha^* = 0$. By Proposition 1, this second alternative cannot occur for $S > 2$. Thus, for $S > 2$, the unique equilibrium in which markets clear has a fraction of agents believing $\hat{E}[Z] < 0$ and shorting

the asset and a fraction of agents believing $\hat{E}[Z] > 0$ and buying the asset. The fractions are such that the aggregate demand for the asset is zero.

We now show that there exists a low enough price such that optimal beliefs always induce the agent to buy the asset. We do this by showing that, for a low enough price, an upper bound on the well-being of an agent that shorts the asset is lower than the well-being of an agent with rational beliefs, who buys the asset. Consider an agent that shorts the asset and consider lower and lower prices for the asset. Since the agent shorts, he must believe $\hat{E}[Z] \leq 0$. As $P \searrow 1 + \varepsilon_1$, $\hat{E}[Z] \leq 0$ implies $\hat{\pi}_1 \nearrow 1$ and $Z_1 \nearrow 0$, so that

$$\begin{aligned} \lim_{P \searrow 1 + \varepsilon_1} \mathcal{W}(\{\hat{\pi}\}, P) &= \lim_{P \searrow 1 + \varepsilon_1} \left\{ \frac{1}{2}\beta \sum_{s=1}^S \hat{\pi}_s u(1 + \alpha^* Z_s) + \frac{1}{2}\beta \sum_{s=1}^S \pi_s u(1 + \alpha^* Z_s) \right\} \\ &= \frac{1}{2}\beta u(1) + \frac{1}{2}\beta \lim_{P \searrow 1 + \varepsilon_1} \sum_{s=1}^S \pi_s u(1 + \alpha^* Z_s) \\ &\leq \frac{1}{2}\beta u(1) + \frac{1}{2}\beta u(1) = \beta u(1) \end{aligned}$$

where the inequality follows from the fact that in the limit the risky asset becomes dominated by the risk-free asset. For $P < \sum_{s=1}^S \pi_s (1 + \varepsilon_s)$, we have $E[Z] > 0$, and so, for an agent with rational beliefs, well-being is

$$\mathcal{W}(\{\pi\}, P) = \beta \sum_{s=1}^S \pi_s u(1 + \alpha^* Z_s)$$

where $\alpha^* > 0$. Since the rational agent chooses α to maximize his objective, α^* yields higher utility than $\alpha = 0$, so

$$\beta \sum_{s=1}^S \pi_s u(1 + \alpha^* Z_s) > \beta u(1)$$

Thus, there is a low enough price such that the beliefs that maximize the well-being function have $\hat{E}[Z] > 0$ and $\alpha^* > 0$. The problem is symmetric, so that there is a completely analogous argument that in the limit as $P \nearrow 1 + \varepsilon_S$, optimal expectations have $\hat{E}[Z] < 0$ and $\alpha^* < 0$.

(ii) The proof for the rational expectations equilibrium is standard and omitted.

For $E[Z^{OE}] = 0$, Proposition 1 directly implies $\alpha^{OE,i} \neq \alpha^{RE} = 0$, because in this case $\alpha^{RE} = \alpha^{RE}(P^{OE}) = \alpha^{RE}(\sum_{s=1}^S \pi_s (1 + \varepsilon_s)) = 0$ where $\alpha^{RE}(P^{OE})$ denotes the portfolio choice of an agent who has beliefs equal to the objective probabilities and faces the optimal

expectations equilibrium price of equity. By market clearing and since agents either short or hold the asset, some agents must be shorting and some agents must be holding the asset, which implies the result. Again by Proposition 1, for $E[Z^{OE}] > 0$, each price-taking agent has $\alpha^{OE,i} > \alpha^{RE}(P^{OE}) > 0$ or $\alpha^{OE,i} < 0$, so $\alpha^{OE,i} \neq 0$, and since $\alpha^{OE,i} \neq \alpha^{RE}(P^{OE})$, $\{\hat{\pi}^{OE,i}\} \neq \{\pi\}$. Again, by market clearing and since agents either short or hold the asset, some agents must be shorting and some agents must be holding the asset, which implies the result. The case of $E[Z^{OE}] > 0$ is analogous.

B.4 Proof of Proposition 4

Part (ii): Subtract $E[c_{t+1}^{OE}|\underline{y}_t]$ from each side of equation (12) with $\tau = 1$

$$\begin{aligned} c_t^{OE} - E[c_{t+1}^{OE}|\underline{y}_t] &= \frac{a}{b} - \frac{\psi_{t+1}}{\psi_t} R \left(\frac{a}{b} - E[c_{t+1}^{OE}|\underline{y}_t] \right) - E[c_{t+1}^{OE}|\underline{y}_t] \\ &= \left(1 - \frac{\psi_{t+1}}{\psi_t} R \right) \left(\frac{a}{b} - E[c_{t+1}^{OE}|\underline{y}_t] \right) \end{aligned}$$

Since the support of the income process does not admit a plan such that $E[c_{t+\tau}^{OE}|\underline{y}_t] > \frac{a}{b}$ for any τ , the second term is positive. The following demonstrates that the first term is positive.

$$\begin{aligned} \frac{\psi_{t+1}}{\psi_t} R &= \frac{\beta^t}{\beta^{t-1}} R \frac{1 + \sum_{\tau=1}^{T-t-1} (\beta^\tau + (\beta\delta)^\tau)}{1 + \sum_{\tau=1}^{T-t} (\beta^\tau + (\beta\delta)^\tau)} \\ &< \beta R \frac{1 + \sum_{\tau=1}^{T-t-1} (\beta^\tau + (\beta\delta)^\tau)}{1 + \sum_{\tau=1}^{T-t-1} (\beta^\tau + (\beta\delta)^\tau) + (\beta^{T-t} + (\beta\delta)^{T-t})} \\ &< 1 \end{aligned}$$

therefore

$$c_t^{OE} - E[c_{t+1}^{OE}|\underline{y}_t] > 0$$

and we have the result.

Part (iii): From the agent's consumption Euler equation

$$\hat{E}[c_{t+1}^{OE}|\underline{y}_t] = c_t^{OE} > E[c_{t+1}^{OE}|\underline{y}_t]$$

where the inequality follows from part (ii).

Part (i): the consumption rule at $t + 1$ is

$$c_{t+1}^{OE} = \frac{1-R^{-1}}{1-R^{-(T-t)}} \left(A_{t+1} + \hat{E} \left[H_{t+1} | \underline{y}_{t+1} \right] \right)$$

where $H_{t+1} = \sum_{\tau=0}^{T-t-1} R^{-\tau} y_{t+\tau+1}$. From part (iii) we have

$$\begin{aligned} \hat{E} \left[c_{t+1}^{OE} | \underline{y}_t \right] &> E \left[c_{t+1}^{OE} | \underline{y}_t \right] \\ \hat{E} \left[\frac{1-R^{-1}}{1-R^{-(T-t)}} \left(A_{t+1} + \hat{E} \left[H_{t+1} | \underline{y}_{t+1} \right] \right) | \underline{y}_t \right] &> E \left[\frac{1-R^{-1}}{1-R^{-(T-t)}} \left(A_{t+1} + \hat{E} \left[H_{t+1} | \underline{y}_{t+1} \right] \right) | \underline{y}_t \right] \end{aligned}$$

and by the law of iterated expectations

$$\hat{E} \left[H_{t+1} | \underline{y}_t \right] > E \left[\hat{E} \left[H_{t+1} | \underline{y}_{t+1} \right] | \underline{y}_t \right].$$

B.5 Proof of Proposition 5

This result obtains because $\frac{\psi_{t+1}}{\psi_t}$ depends only on the number of periods until T , $T - t$:

$$\frac{\psi_{t+1}}{\psi_t} = \beta \frac{1 + \sum_{\tau=1}^{T-t-1} (\beta^\tau + (\beta\delta)^\tau)}{1 + \sum_{\tau=1}^{T-t} (\beta^\tau + (\beta\delta)^\tau)}$$

Thus, optimal beliefs from the perspective of any time period imply the same rationally expected percent change in the profile of marginal utility between any two periods. Since the budget constraint determines the level of the profile and all plans exhaust the resources, the levels are necessarily the same. Since income is perceived as certain, if that income actually occurs, beliefs are bound by Bayes' rule. In this case, there is no change in beliefs about future incomes, and the subjective expectation of consumption coincides with actual consumption.

References

- AKERLOF, G., AND W. T. DICKENS (1982): “The Economic Consequences of Cognitive Dissonance,” *American Economic Review*, 72(3), 307–319.
- ALPERT, M., AND H. RAIFFA (1982): “A Progress Report on the Training of Probability Assessors,” in *Judgement Under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, pp. 294–305. Cambridge University Press, Cambridge.
- ATTANASIO, O. (1999): “Consumption,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 11, pp. 741–812. Elsevier Press, Amsterdam.
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *Annals of Statistics*, 4, 1236–1239.
- BARGH, J. A., AND T. L. CHARTRAND (1999): “The Unbearable Automaticity of Being,” *American Psychologist*, 54(7), 462–479.
- BARSKY, R., T. JUSTER, M. KIMBALL, AND M. SHAPIRO (1997): “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey,” *Quarterly Journal of Economics*, 112(2), 537–580.
- BÉNABOU, R., AND J. TIROLE (2002): “Self-confidence and Personal Motivation,” *Quarterly Journal of Economics*, 117(3), 871–915.
- (2004): “Willpower and Personal Rules,” *Journal of Political Economy*, 112(4), 848–86.
- BERNHEIM, B. D., AND R. THOMADSEN (2003): “Memory and Anticipation,” *Mimeo, Stanford University*.
- BRUNNERMEIER, M. K., AND J. A. PARKER (2002): “Optimal Expectations,” *Woodrow Wilson School Discussion Papers in Economics*, 221.
- BUEHLER, R., D. GRIFFIN, AND M. ROSS (1994): “Exploring the “Planning Fallacy”: Why People Underestimate Their Task Completion Times,” *Journal of Personality and Social Psychology*, 67(3), 366–381.
- CAPLIN, A. J., AND J. LEAHY (2000): “The Social Discount Rate,” *NBER Working Paper*, 7983.
- (2001): “Psychological Expected Utility Theory and Anticipatory Feelings,” *Quarterly Journal of Economics*, 106(1), 55–80.
- (2004): “The Supply of Information by a Concerned Expert,” *Economic Journal*, 114(497), 487–505.

- CARRILLO, J. D., AND T. MARIOTTI (2000): “Strategic Ignorance as a Self-Disciplining Device,” *Review of Economic Studies*, 67, 529–544.
- CHEN, J., H. HONG, AND J. C. STEIN (2001): “Forecasting Crashes: Trading Volume, Past Returns and Conditional Skewness in Stock Prices,” *Journal of Financial Economics*, 61, 345–381.
- COHEN, S., W. J. DOYLE, R. B. TURNER, C. M. ALPER, AND D. P. SKONER (2003): “Emotional Style and Susceptibility to the Common Cold,” *Psychosomatic Medicine*, 65, 652–657.
- DELONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. J. WALDMANN (1990): “Noise Trader Risk in Financial Markets,” *Journal of Political Economy*, 98(4), 703–738.
- ELIAZ, K., AND R. SPIEGLER (2003): “Anticipatory Feelings and Choices of Information Sources,” *Mimeo, University of New York*.
- EYSTER, E. (2002): “Rationalizing the Past: A Taste for Consistency,” *Mimeo, Oxford University*.
- GEANAKOPOLOS, J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, 1, 60–79.
- GOLEC, J., AND M. TAMARKIN (1998): “Gamblers Love Skewness, Not Risk, at the Horse Track,” *Journal of Political Economy*, 106, 205–255.
- GOURINCHAS, P.-O., AND J. A. PARKER (2002): “Consumption Over the Lifecycle,” *Econometrica*, 70(1), 47–89.
- HARBAUGH, R. (2002): “Skill Reputation, Prospect Theory, and Regret Theory,” *Mimeo, Claremont Colleges*.
- KAHNEMAN, D. (2000): “Experienced Utility and Objective Happiness: A Moment-Based Approach,” in *Choices, Values and Frames*, ed. by D. Kahneman, and A. Tversky, pp. 637–692. Cambridge University Press, New York.
- KAHNEMAN, D., P. P. WAKKER, AND R. SARIN (1997): “Back To Bentham? Explorations of Experienced Utility,” *Quarterly Journal of Economics*, 112(2), 375–405.
- KIECOLT-GLASER, J. K., L. MCGUIRE, T. F. ROBLES, AND R. GLASER (2002): “Psychoneuroimmunology and Psychosomatic Medicine: Back to the Future,” *Psychosomatic Medicine*, 64, 15–28.
- LANDIER, A. (2000): “Wishful Thinking: A Model of Optimal Reality Denial,” *Mimeo, MIT*.

- LEWIS, K. K. (1999): “Trying to Explain Home Bias in Equity and Consumption,” *Journal of Economic Literature*, 37(2), 571–608.
- LICHTENSTEIN, S., B. FISCHHOFF, AND L. D. PHILLIPS (1982): “Calibration of Probabilities: The State of the Art to 1980,” in *Judgement Under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, pp. 306–334. Cambridge University Press, Cambridge.
- LOEWENSTEIN, G. (1987): “Anticipation and the Valuation of Delayed Consumption,” *Economic Journal*, 47(387), 666–684.
- LUCAS, JR., R. E. (1976): “Econometric Policy Evaluation: A Critique,” in *The Phillips Curve and Labor Markets*, ed. by K. Brunner, and A. H. Meltzer, vol. 1 of *Carnegie Rochester Conference Series on Public Policy*, pp. 19–46. North-Holland, Amsterdam.
- MULLAINATHAN, S. (2002): “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 117(3), 735–774.
- MUTH, J. F. (1960): “Optimal Properties of Exponentially Weighted Forecasts,” *Journal of the American Statistical Association*, 55, 229–305.
- (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29(3), 315–335.
- MYERSON, R. B. (1986): “Axiomatic Foundations of Bayesian Decision Theory,” *Northwestern University, MEDS Discussion Paper*, 671.
- NERLOVE, M. (1958): “Adaptive Expectations and Cobweb Phenomena,” *Quarterly Journal of Economics*, 72, 227–240.
- PICCIONE, M., AND A. RUBINSTEIN (1997): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20, 3–24.
- POTERBA, J. M. (2003): “Employer Stock and 401(k) Plans,” *American Economic Review (Papers and Proceedings)*, 93(2), 398–404.
- RABIN, M., AND J. L. SCHRAG (1999): “First Impressions Matter: A Model of Confirmatory Bias,” *Quarterly Journal of Economics*, 114(1), 37–82.
- SAVAGE, L. J. (1954): *The Foundation of Statistics*. Wiley, New York.
- WEINSTEIN, N. D. (1980): “Unrealistic Optimism About Future Life Events,” *Journal of Personality and Social Psychology*, 39(5), 806–820.
- YARIV, L. (2001): “Believe and Let Believe: Axiomatic Foundations for Belief Dependent Utility,” *Yale University, Cowles Foundation Discussion Paper*, 1344.

——— (2002): “I’ll See It When I Believe It – A Simple Model of Cognitive Consistency,”
Mimeo, UCLA.