

# Risk Attitude Optimization and Heterogeneous Stock Market Participation\*

Todd Sarver<sup>†</sup>

October 20, 2017

## Abstract

This paper studies equilibrium portfolio choice and asset returns using a new model of recursive preferences called optimal risk attitude utility. Our model is an extension of recursive expected utility that allows an individual to optimally select her risk aversion parameter in response to the uncertainty that she faces. Choosing a lower level of risk aversion comes at a cognitive cost, and therefore is only undertaken in response to sufficiently large risk exposure. In addition to separating risk aversion from the elasticity of intertemporal substitution, our model can also separate risk attitudes toward small and large risks. We solve the dynamic stochastic general equilibrium of a calibrated economy and show that optimal risk attitude utility can provide a partial resolution to both the stock market participation puzzle and the equity premium puzzle. Our model generates a moderate premium for equity as well as endogenous heterogeneity in risk exposure, with one segment of the population holding minimal risk while the other holds a disproportionate share of aggregate risk, even when consumers have identical preferences and even among wealthy households. Using simple binary gambles as a rationality check for our model, we demonstrate that the preference parameters used in our calibration exhibit descriptively accurate levels of risk aversion for gambles ranging from one hundred dollars up to ten percent of wealth.

KEYWORDS: optimal risk attitude, stock market participation puzzle, equity premium puzzle

---

\*I thank David Ahn, Tim Bollerslev, Cosmin Ilut, Jonathan Parker, Wolfgang Pesendorfer, Philipp Sadowski, and seminar participants at University of Pennsylvania for helpful comments and discussions.

<sup>†</sup>Duke University, Department of Economics, 213 Social Sciences/Box 90097, Durham, NC 27708.  
Email: todd.sarver@duke.edu.

# 1 Introduction

As documented by [Mankiw and Zeldes \(1991\)](#), [Haliassos and Bertaut \(1995\)](#), [Heaton and Lucas \(2000\)](#), and others, a significant number of households do not participate in equity markets. Even with the recent upward trend in stock market investing through retirement savings and mutual funds, as of 2007 only 51.5% of US households have any direct or indirect holdings of stock, and the average share of financial assets held in equity is just 52.7% ([Guiso and Sodini \(2013\)](#)). In addition, although participation is positively correlated with wealth, a nontrivial fraction of wealthy households hold little or no public or private equity.<sup>1</sup> Indeed, as noted by [Campbell \(2006, page 1564\)](#), “Limited participation among the wealthy poses a significant challenge to financial theory and is one of the main stylized facts of household finance.”

These empirical observations are often referred to as the stock market *participation puzzle*. Much like the equity premium puzzle ([Mehra and Prescott \(1985\)](#); [Kocherlakota \(1996\)](#)), the participation puzzle is quantitative in nature. The equity premium puzzle refers to the inability of the expected-utility model to match the observed equity premium of roughly 6% without imposing an implausibly high coefficient of risk aversion, given the low historic risk in the stock market as measured by its co-movement with *aggregate* consumption growth. Similarly, the participation puzzle refers to the inability of expected utility to generate nonparticipation using a reasonable coefficient of risk aversion, given the even lower co-movement of the stock market with the consumption of *nonparticipating* households.

We propose a partial resolution to both the stock market participation puzzle and the equity premium puzzle by using *optimal risk attitude (ORA)* utility, a new model of recursive utility that is developed and studied axiomatically in the companion paper [Sarver \(2017\)](#). ORA utility is an extension of dynamic expected utility in which an individual is able to optimize over her risk attitude in response to the uncertainty that she faces. Formally, the special case of ORA preferences considered in this paper admits the following value function for any random consumption stream  $c^t = (c_t, c_{t+1}, \dots)$ :

$$\mathcal{V}(c^t) = (1 - \beta) \log(c_t) + \beta \sup_{\theta \in \Theta} \left\{ -\frac{1}{\theta} \log(\mathbb{E}_t[\exp(-\theta \mathcal{V}(c^{t+1}))]) - \tau(\theta) \right\}.$$

The elements of the utility representation are a discount factor  $\beta$ , a set of risk-aversion parameters  $\Theta \subset \mathbb{R}_+$ , and a cost function  $\tau$  that satisfies  $\inf_{\theta \in \Theta} \tau(\theta) = 0$ . These preferences have a simple interpretation in terms of costly psychological preparation: An individual may mentally prepare herself for different levels of risk by optimizing over the

---

<sup>1</sup>[Heaton and Lucas \(2000\)](#) showed that while private business assets substitute for public equity for some households, around 10% of wealthy households hold neither.

parameter  $\theta$ , where lower values of this parameter decrease her sensitivity to risk but may also come at a higher cognitive cost  $\tau(\theta)$ .

In the special case of a single parameter,  $\Theta = \{\theta\}$  and  $\tau(\theta) = 0$ , this model reduces to recursive expected utility in the sense of [Kreps and Porteus \(1978\)](#) and [Epstein and Zin \(1989\)](#), which we will refer to as Epstein-Zin-Kreps-Porteus (EZKP) utility. As emphasized by [Epstein and Zin \(1989, 1991\)](#) and [Weil \(1989, 1990\)](#), EZKP utility permits a separation between risk aversion and the elasticity of intertemporal substitution that is not possible for standard time-separable expected utility. This separation is useful, for example, in allowing the model to match both the observed high equity premium and the low risk-free rate.<sup>2</sup>


Despite its benefits, one important limitation of EZKP utility is its inability to separate attitudes toward small and large risks: It cannot generate realistically high levels of risk aversion for small- or medium-stakes gambles without also producing excessively high levels of risk aversion for large-stake gambles.<sup>3</sup> This tight linkage between risk attitudes across different scales of risk is at the heart of both the equity premium puzzle and the participation puzzle. For example, it is known that EZKP utility can generate an equity premium of around 6% by taking a coefficient of relative risk aversion of roughly 18 ([Kocherlakota \(1996\)](#)). However, [Mehra and Prescott \(1985\)](#), [Kocherlakota \(1996\)](#), [Lucas \(2003\)](#), and others have argued that fitting the data using expected utility with a CRRA greater than 10 is unreasonable, as the implied risk aversion for large-scale idiosyncratic risks (e.g., wage premiums for occupations with high earnings risk) is absurdly high. Thus imposing the level of risk aversion needed to generate a high equity premium in response to the moderate risk associated with aggregate consumption growth requires excessively high risk aversion in other domains.

Optimal risk attitude utility addresses precisely this linkage between risk aversion for small and large gambles. For any fixed  $\theta$ , ORA utility evaluates uncertain consumption according to expected utility with a fixed coefficient of risk aversion. However, since  $\theta$  can vary with the risk being faced, different *local* risk attitudes may be exhibited for different risks. For example, suppose  $\Theta = \{\theta_L, \theta_H\}$  for  $\theta_H > \theta_L$  and  $0 = \tau(\theta_H) < \tau(\theta_L)$ . For small gambles it will be optimal to choose  $\theta_H$  in order to avoid the utility cost  $\tau(\theta_L) > 0$ ; when faced with larger risks, a consumer may instead find it optimal to select  $\theta_L$  to decrease her sensitivity to this more uncertain outcome. More generally, the optimization over risk

---

<sup>2</sup>The exponential form of EZKP utility corresponding to  $\Theta = \{\theta\}$  in our model has been used in a number of macroeconomic applications (e.g., [Hansen, Sargent, and Tallarini \(1999\)](#); [Tallarini \(2000\)](#)). It has also been interpreted in terms of robustness to model uncertainty, as its equivalence with the multiplier preferences of [Hansen and Sargent \(2001\)](#) has been established in a variety of settings (e.g., [Skiadas \(2003\)](#); [Maenhout \(2004\)](#); [Strzalecki \(2011\)](#)). For a detailed discussion of this reinterpretation in the context of the equity premium puzzle, see [Barillas, Hansen, and Sargent \(2009\)](#).

<sup>3</sup>The calibration result of [Rabin \(2000\)](#), known as the Rabin paradox, provides a convincing illustration of this property of expected utility using simple binary gambles.



**Figure 1.** Indifference curves and endogenous heterogeneity in general equilibrium illustrated for a static model with two states. Type  $\theta_L$  consumers choose allocation  $c_{\theta_L}$  and type  $\theta_H$  consumers choose allocation  $c_{\theta_H}$ . The fraction  $\alpha$  of consumers selecting type  $\theta_L$  is determined by the market-clearing condition:  $e = \alpha c_{\theta_L} + (1 - \alpha) c_{\theta_H}$ .

attitudes in our model implies that ORA utility will violate the independence axiom of expected utility. Sarver (2017) shows that these preferences will instead satisfy a weaker axiom called *mixture aversion* that is connected to the certainty effect (Allais paradox), probabilistic insurance, and other experimental evidence concerning choice under risk.

In this paper, we consider an economy with a continuum of agents with identical ORA preferences that are homothetic in wealth. We find that heterogeneity in risk exposure can arise endogenously in the dynamic general equilibrium of this economy. The intuition behind our results is relatively straightforward and is illustrated in Figure 1 for a stylized static model with two states ( $z^l$  and  $z^h$ ) where consumers have identical endowments  $e$ : When agents can optimize over their risk attitudes subject to some mental cost, they may endogenously sort into different types in equilibrium. One segment of the population will choose a risk attitude  $\theta_L$  that is less sensitive to gains and losses and therefore will hold greater consumption risk; the other will choose a risk attitude  $\theta_H$  that provides higher utility for low-risk allocations ( $\tau(\theta_H) < \tau(\theta_L)$ ), but is more sensitive to losses ( $\theta_H > \theta_L$ ), and will hold very little consumption risk. This equilibrium allocation with heterogeneous risk exposure is a Pareto improvement of the homogeneous allocation where each agent consumes her endowment. Moreover, heterogeneity is robust in our model: This is not just one possible equilibrium, but rather is the unique equilibrium of our economy for a range of parameter values.

It is worth emphasizing that although equilibrium in our model will involve significant heterogeneity in risk holding, with one segment of the population bearing minimal risk, equilibrium in our baseline model will not involve the complete absence of risk for any segment of the population. Thus in the simple case where all uncertainty about consumption is driven by investment risk, our model will not imply *complete* nonparticipation. However, as we discuss in Section 6.1, the utility difference between participation and nonparticipation becomes negligible in our model. As a result, introducing a minimal friction in the form of a small stock market participation cost will be sufficient to obtain complete nonparticipation by a segment of the population.<sup>4</sup> In contrast, as we discuss in Section 2, nonparticipation by wealthy households cannot be explained using reasonable values for the participation cost in a standard expected-utility framework.

In the dynamic general equilibrium of our model, inducing a fraction of the population to hold a disproportionate share of aggregate risk requires significant compensation in the form of a larger expected return. In this way, our model will generate both heterogeneous participation and a large equity premium. This mechanism accords well with a large body of empirical evidence: It has been shown that the consumption of stockholders is more volatile and more correlated with stock market returns than that of nonstockholders, and hence the equity premium becomes less of a puzzle when attention is restricted to consumers who invest in the stock market (see Mankiw and Zeldes (1991); Attanasio, Banks, and Tanner (2002); Brav, Constantinides, and Geczy (2002); Vissing-Jørgensen (2002); Vissing-Jørgensen and Attanasio (2003)).

At the same time, the aforementioned papers observe that explaining the nonparticipation of the remaining households poses a deeper puzzle for expected utility, as the level of risk aversion required to rationalize their choices would be even higher than estimates based on aggregate consumption. However, for ORA utility, the high imputed coefficients of risk aversion for households holding minimal consumption risk in equilibrium do not evoke the same concerns that are discussed, for example, in Lucas (2003) and Campbell (2003) in the context of expected utility. In our model, nonparticipants have high levels of *local* risk aversion, but do not exhibit overly high levels of *global* risk aversion (recall that their preferences are ex-ante identical to those of participants). In Section 5, we use preferences over binary gambles to paint a clear picture of the overall risk attitudes implied by different parameter values in our model.

The remainder of the paper is organized as follows. In Section 2, we discuss other explanations that have been proposed to address the participation puzzle and describe how they relate to our approach. In Section 3, we formally describe our economy with a continuum of consumers with identical ORA preferences that are homothetic in wealth.

---

<sup>4</sup>Alternatively, incorporating kinked indifference curves due to first-order risk aversion or ambiguity aversion or including background risk with sufficient correlation with stock market risk could also lead to complete nonparticipation in our model.

In Section 4, we illustrate informally how heterogeneity in risk exposure can arise endogenously in this economy, with some agents holding very little risk in equilibrium— independent of their wealth. In Section 5, we provide a calibration of the dynamic general equilibrium of this economy that verifies these informal observations and quantifies the level of participation and the resulting equity premium. Section 6 concludes with a discussion of possible extensions of the model.

## 2 Existing Explanations of Limited Participation

The stock market participation puzzle has received significant attention in recent years, and a number of potential explanations have been proposed. In this section, we summarize the successes (and shortcomings) of existing models in addressing these facts, and discuss the similarities and differences from our approach.

EXPLANATION 1: Participation costs.

Participation costs, in the form of monetary expenses associated with investing or informational costs associated with choosing an optimal portfolio of stocks, can provide a partial resolution of the participation puzzle. A number of studies have found that plausible values of entry costs and ongoing participation costs can rationalize the non-participation decision of many households (e.g., [Vissing-Jørgensen \(2003\)](#); [Gomes and Michaelides \(2005\)](#)). Another important benefit of participation cost models is that they predict that participation increases with wealth. However, these models still fail to explain the lack of participation or minimal risk exposure of some wealthy households, as their benefit from equity investing would dwarf any reasonable value for the participation cost ([Vissing-Jørgensen \(2003\)](#); [Briggs, Cesarini, Lindqvist, Östling \(2015\)](#)).

We view our model as complementary to the participation-cost approach. Since ORA utility is homothetic in wealth, our predictions regarding heterogeneous participation in no way rely on wealth effects. In Section 6.1, we discuss how an extension of our model that combines ORA utility with modest participation costs can explain why risk exposure covaries positively with wealth, yet even slight heterogeneity in costs or preferences in the population would lead some wealthy households to either not participate in the stock market or hold very conservative portfolios.

EXPLANATION 2: First-order risk aversion.

Other preference-based models have been used to address the participation decisions of wealthy households. These models invoke first-order risk aversion ([Segal and Spivak](#)

(1990)) as an explanation for why some households would avoid an actuarially favorable investment opportunity such as the stock market. For example, [Ang, Bekaert, and Liu \(2005\)](#) used the disappointment aversion model of [Gul \(1991\)](#) to study a dynamic asset allocation problem and determined the critical values of the disappointment aversion parameter that lead to nonparticipation. However, [Barberis, Huang, and Thaler \(2006\)](#) observed that the presence of background risk (e.g., uninsurable idiosyncratic income risk) makes it difficult for models of first-order risk aversion to explain nonparticipation in the stock market using reasonable parameter values.<sup>5</sup> Their suggested remedy was to assume loss aversion together with narrow framing of portfolio risk, meaning that gains or losses in the stock market are evaluated separately from overall consumption risk. In [Section 6.2](#), we discuss why our model is unlikely to be subject to their critique and thus narrow framing is not needed to generate low participation levels. To the contrary, we will show that existing results about the impact of background risk on the risk attitudes of expected-utility maximizers imply that moderate background risk would only serve to *decrease* the optimal level of investment of households in our model.

### EXPLANATION 3: Ambiguity aversion.

Ambiguity aversion has also been proposed as an explanation for nonparticipation in the stock market. It is well known from the work of [Dow and Werlang \(1992\)](#) and [Epstein and Wang \(1994\)](#) that ambiguity aversion can lead to portfolio inertia at the risk-free portfolio, and therefore heterogeneity in perceived ambiguity can generate nonparticipation by some agents. More recent work by [Epstein and Schneider \(2007\)](#) has explored how learning under ambiguity can influence stock market participation, and they showed that this mechanism may explain part of the increase in participation rates in recent years.<sup>6</sup> Ambiguity aversion and optimal risk attitude utility are in many ways complementary—as was the case with participation costs—in that each is ideally suited to address different aspects of household behavior. For example, ambiguity aversion provides a sensible rationale for underdiversification and the home bias (e.g., [Epstein and Miao \(2003\)](#); [Boyle, Garlappi, Uppal, and Wang \(2012\)](#)), neither of which has an obvious connection to our model. On the other hand, ORA utility can generate low levels

---

<sup>5</sup>[Barberis, Huang, and Thaler \(2006\)](#) showed that introducing background risk decreases the aversion to exposure to stock market risk for the disappointment aversion model, to the point where the disappointment aversion parameters required for nonparticipation in the stock market would also imply rejection of a 50-50 gamble with a loss of \$10,000 and a gain of \$20,000,000 at a wealth level of \$100,000. They argued that their negative result extends broadly to models that rely on first-order risk aversion to obtain nonparticipation. [Safra and Segal \(2008\)](#) made the related observation that Rabin’s paradox extends to a broad class of non-expected-utility preferences when sufficient independent background risk is introduced.

<sup>6</sup>Another likely explanation for rising participation rates is a decrease in informational and monetary costs to investing (e.g., see the discussion in [Guiso and Sodini \(2013, page 1454\)](#)).

of stock market participation and low portfolio weights on stock for many participating households even in familiar environments or after long histories, where ambiguity would arguably play a less significant role in investment decisions.

Due to the complexity of dynamic models with heterogeneous preferences, the extant literature on preference-based explanations of the participation puzzle has been restricted almost exclusively to partial equilibrium analysis, taking the data-generating process for asset returns as given and solving for the participation and asset allocation decisions of a single agent. The few papers that have conducted general equilibrium analysis (typically within a one or two period framework) have reported somewhat mixed results. For example, [Cao, Wang, and Zhang \(2005\)](#) showed that increasing the ambiguity dispersion among investors leads simultaneously to a decrease in the participation rate and a decrease in the equity premium.<sup>7</sup> [Chapman and Polkovnichenko \(2009\)](#) observed similar implications for increases in the dispersion of risk aversion parameters for a variety of non-expected-utility preferences (including disappointment aversion and rank-dependent utility). These results suggest that more research is needed to determine the suitability of these models for jointly explaining the participation rate and asset returns.

This is perhaps the most important point of departure of our model from the previous literature: It can generate equilibrium heterogeneity in participation decisions even when agents have identical and homothetic preferences. Although our model will not permit representative agent analysis, these conditions imply that equilibrium analysis does not require tracking the distribution of wealth across agents. Consequently, the model is tractable enough for us to conduct a complete dynamic general equilibrium analysis of both market prices and the distribution of participation decisions, taking only the dividend and consumption processes as primitives. This makes it possible to easily evaluate the performance of ORA utility as a driver of both nonparticipation and stock returns.

## 3 Model

### 3.1 Economy

Aggregate consumption growth depends on a state variable  $z_t \in Z$  that is observed by each individual in the economy at the start of period  $t$ , where  $Z$  is finite. We assume that the state is i.i.d. across time, and is distributed according to a probability measure  $P$ . Thus the model precludes intertemporal correlation between current and future states. This assumption is certainly restrictive, and is made in order to simplify the analysis and

---

<sup>7</sup>The intuition behind their result is that greater dispersion in ambiguity implies there is a subset of the population that becomes much more tolerant of ambiguous payoffs and is therefore willing to invest more heavily in stock even at a lower premium.



focus attention on the endogenous heterogeneity in risk bearing that arises in equilibrium. A natural next step for future study is to incorporate persistence of shocks and other types of correlation across time into this analysis.

Denote the aggregate consumption endowment by  $e_t$ , and let  $\lambda_{t+1}$  denote aggregate consumption growth between period  $t$  and  $t + 1$ :  $e_{t+1} = \lambda_{t+1}e_t$ . Consumption growth is determined by the state, i.e,  $\lambda_{t+1} = \lambda(z_{t+1})$  for some function  $\lambda$ . We assume a continuum of consumers, with the consumption of consumer  $i \in [0, 1]$  at time  $t$  denoted by  $c_{i,t}$ . The feasibility constraint in the economy requires that

$$\int_0^1 c_{i,t} di = e_t, \quad \forall t.$$

The distribution of ownership of the endowment allocation is not important for any of the results that follow, as long as it is absolutely continuous with respect to the Lebesgue measure on the unit interval of consumers.

Suppose that markets are complete and there exists a pricing kernel  $M_{t,t+1}$  such that the period  $t$  price of any asset paying  $x_{t+1}$  in period  $t + 1$  is  $p_t = \mathbb{E}_t[M_{t,t+1}x_{t+1}]$ . Given the independence of consumption growth across time, one might conjecture that the equilibrium pricing kernel is stationary and depends only on the period  $t + 1$  state:  $M_{t,t+1} = M(z_{t+1})$  for some function  $M$ . We will verify that this is indeed the case in equilibrium. Under these assumptions, the time  $t$  budget constraint for a consumer with wealth  $w_t$  is then

$$c_t + \mathbb{E}[Mw_{t+1}] = w_t, \tag{1}$$

where  $c_t \in \mathbb{R}_+$  and  $w_{t+1} \in \mathbb{R}_+^Z$ .

### 3.2 Preferences

We model the preferences of the consumers in this economy using a special case of the *optimal risk attitude (ORA)* recursive utility studied in Sarver (2017). In our utility representation, individuals can optimize over their risk-aversion parameter  $\theta$  at some unobserved (psychological) cost. One interpretation of this model is that individuals can “desensitize” themselves to risk, but doing so requires some mental effort or psychological discomfort. Formally, the value function given wealth  $w_t$  and state prices  $M$  is

$$\mathcal{V}(w_t; M) = \max_{\substack{c_t \in \mathbb{R}_+ \\ w_{t+1} \in \mathbb{R}_+^Z}} \left\{ (1 - \beta) \log(c_t) + \beta \mathcal{R}(\mathcal{V}(w_{t+1}; M)) \right\}, \tag{2}$$

where the maximization is subject to the budget constraint in Equation (1), and  $\mathcal{R}$  is an operator that maps any random continuation value  $\mathcal{V}(w_{t+1}; M)$  into a *risk-adjusted*

*continuation value.* The risk preferences of the individual are fully captured by this operator, which takes the following functional form in our model:

$$\mathcal{R}(\mathcal{V}(w_{t+1}; M)) = \sup_{\theta \in \Theta} \left\{ -\frac{1}{\theta} \log \left( \mathbb{E} \left[ \exp(-\theta \mathcal{V}(w_{t+1}; M)) \right] \right) - \tau(\theta) \right\}. \quad (3)$$

In this equation, risk preferences are specified by two objects:  $\Theta \subset \mathbb{R}_+$  is a set of feasible risk-aversion parameters and  $\tau : \Theta \rightarrow \mathbb{R}$  is a cost function that satisfies  $\inf_{\theta \in \Theta} \tau(\theta) = 0$ .<sup>8</sup>

The objects  $\Theta$  and  $\tau$  are not directly observable. Rather, they are parameters of the utility function, much like a coefficient of risk aversion in the standard expected-utility model, and they must therefore instead be inferred from an individual's choices. For example, Sarver (2017) shows that an individual with a lower cost function will be less risk averse. In Section 5, we will use preferences over binary gambles to illustrate the overall risk attitudes associated with different parameter values.

As noted in the introduction, our model reduces to recursive expected utility in the sense of Kreps and Porteus (1978) and Epstein and Zin (1989) if  $\Theta = \{\theta\}$  and  $\tau(\theta) = 0$ . We will refer to this standard model as Epstein-Zin-Kreps-Porteus (EZKP) utility. The equilibrium analysis in Section 5 will include two specifications of EZKP utility for comparison. Our results will show that expanding this classic model to permit multiple values of  $\theta$  leads to several significant differences from EZKP preferences: First, optimizing over risk attitudes permits individuals to be very averse to small gambles while at the same time exhibiting reasonable risk attitudes toward larger gambles. We will see that this feature helps the model to address puzzles such as the Rabin paradox and the equity premium puzzle, since the empirical challenge underlying both is essentially that of simultaneously matching observed levels of risk aversion for both large and small risks using a single model. Second, equilibrium in this economy often involves heterogeneity in the choice of  $\theta$  and, as a result, significant cross-sectional variation in the consumption risk held by consumers, even when they are ex-ante identical.

Standard techniques can be applied to prove the existence and uniqueness of the value function for the consumer's problem.

**Proposition 1** *For any pricing kernel  $M : Z \rightarrow \mathbb{R}_{++}$ , there exists a unique value function  $\mathcal{V}$  for the problem described in Equations (1), (2), and (3). This value function takes the form*

$$\mathcal{V}(w_t; M) = \Lambda(M) + \log(w_t), \quad (4)$$

for a constant  $\Lambda(M)$ .

---

<sup>8</sup>Setting the minimum cost equal to zero is a convenient normalization and is without loss of generality whenever the cost function is bounded below. Adding or subtracting a constant from the function  $\tau$  shifts the value function by a scalar, but does not alter the consumer's preferences.

The proof of this result is contained in the appendix, where we also describe the specific formula for  $\Lambda(M)$  and the optimal consumption and future wealth conditional on each state  $z$  and each possible choice of parameter  $\theta$ .

We conclude this section with a brief comment about additional generality that could be incorporated into the model. The most general specification of optimal risk attitude utility that is used to represent the class of mixture-averse preferences in Sarver (2017) is broad enough to also incorporate first-order risk aversion, i.e., kinks in preferences at the certainty line. First-order risk aversion can lead consumers to fully insure even at actuarially unfair rates, or completely abstain from investing in an actuarially-favorable asset. We eschew first-order risk aversion and instead focus on the risk preferences described in Equation (3) for several reasons: First, we want to make the model as standard as possible by taking only a slight (one-parameter) deviation from recursive expected utility. Second, despite its apparent usefulness, our results will show that first-order risk aversion is not needed to generate a fraction of the population that holds minimal risk exposure. Third, from an analytical perspective, our utility function is smooth and therefore has the convenience of permitting techniques based on differentiation. Finally, our utility function allows the model to make realistic predictions about the impact of incorporating uninsurable idiosyncratic income risk (see Section 6.2).

## 4 Illustration of Equilibrium

In this section, we provide an informal description of the three possible types of equilibrium that can arise in this model. For ease of illustration, the discussion will initially focus on equilibrium in a static (one-period) model. We then proceed to describe how equilibrium in our infinite-horizon model has a similar structure. Formal statements of the equilibrium conditions for the three cases are relegated to Appendix A.2, which also contains a summary of the numerical procedure that will be used in Section 5.

### 4.1 Benchmark: Static Model

For ease of illustration, we first consider a simple model with a single period, two states, and two types. This section will therefore focus on the following simple specification:

$$Z = \{z^l, z^h\}, \quad e(z^l) < e(z^h),$$

$$\Theta = \{\theta_L, \theta_H\}, \quad \theta_H > \theta_L, \quad \tau(\theta) = \begin{cases} \tau_H = 0 & \text{if } \theta = \theta_H \\ \tau_L > 0 & \text{if } \theta = \theta_L. \end{cases}$$

We have dropped time subscripts for the one-period analysis in this section. Assume each of the continuum of consumers has equal (unit) ownership of the endowment (this assumption is made for ease of exposition but is not imposed in the analysis in the following section). Thus a consumption allocation  $c = (c(z^l), c(z^h))$  is affordable given the pricing kernel  $M$  if

$$\mathbb{E}[Mc] = \mathbb{E}[Me].$$

Assume each consumer evaluates a random consumption allocation  $c$  by applying the risk-adjustment operator in Equation (3) to the logarithm of  $c(z)$ :<sup>9</sup>

$$\mathcal{R}(\log(c)) = \max_{\theta \in \{\theta_L, \theta_H\}} \left\{ \log \left( \mathbb{E}[c^{-\theta}]^{-\frac{1}{\theta}} \right) - \tau(\theta) \right\}.$$

Thus, for fixed  $\theta$ , each consumer evaluates uncertain consumption using a CRRA certainty equivalent with a coefficient of relative risk aversion of  $\theta + 1$ . However, since  $\theta$  can vary with the risk being faced, this preference will violate the expected-utility axioms. Intuitively, for small gambles it will be optimal to choose  $\theta_H$  in order to avoid the utility cost  $\tau_L > 0$ ; when faced with larger risks, a consumer may then find it optimal to select  $\theta_L$  to decrease her sensitivity to this more uncertain outcome.


Figures 1 and 2 illustrate the indifference curves of a consumer with the risk preferences described above. Note that  $\theta_H$  is optimal for allocations near the certainty line, whereas  $\theta_L$  is optimal for allocations where the level of consumption differs greatly between states. These figures also illustrate the equilibrium allocations and budget lines for three possible cases. In cases 1 and 2, which are illustrated in Figure 2, all consumers choose the same type in equilibrium ( $\theta_L$  and  $\theta_H$ , respectively) and consume their endowment:  $c(z) = e(z)$  for  $z \in \{z^l, z^h\}$ . In these first two cases, the equilibrium allocation and prices are the same for our economy with a continuum of consumers as they would be for an economy with a single representative agent with the same preferences.

The distinctive part of our analysis arises in case 3, which was illustrated in Figure 1 in the introduction. In this case, equilibrium necessarily involves heterogeneous types, with some consumers selecting  $\theta_H$  and others  $\theta_L$ . As is evident from the figure, preferences are not quasiconcave in the consumption allocation, and therefore equilibrium may not exist in a representative agent economy (non-existence would be an issue precisely in case 3). However, for our economy with a continuum of consumers, the theorem of Aumann (1966) ensures the existence of an equilibrium.<sup>10</sup> The crux of his theorem is that the average

---

<sup>9</sup>To make this analysis as representative as possible of the infinite-horizon model, we treat consumption just as we would treat future wealth in the general model. Given the log form of the value function in Equation (4), we therefore apply the risk-adjustment operator  $\mathcal{R}$  to  $\log(c)$  rather than to  $c$ .

<sup>10</sup>A related result for a large but finite economy in which consumers' preferences are permitted to violate quasiconcavity can be found in Starr (1969), who showed that divergence from equilibrium shrinks with the number of consumers.



**Figure 2.** Two possible cases of homogeneous equilibrium illustrated for a static model with two states. In both cases, equilibrium consumption is identical for each consumer and proportional to the aggregate endowment. Figure 1 in the introduction illustrates the third case of heterogeneous equilibrium consumption.

(i.e., integral) of the (possibly non-convex) upper-contour sets of individual preferences is convex. The economic interpretation of this mathematical condition is central to our results: Rather than every consumer keeping their endowment, there may be welfare improvements associated with heterogeneous allocations that lead to the same average aggregate consumption. Equilibrium in case 3 takes this form, with consumers who select type  $\theta_L$  choosing allocation  $c_{\theta_L}$  and consumers who select type  $\theta_H$  choosing allocation  $c_{\theta_H}$ . Note that these two allocations give the same utility, which is strictly higher than the utility from consuming the endowment  $e$ . The fraction of consumers selecting each of these types and allocations is determined by the market-clearing condition:  $e = \alpha c_{\theta_L} + (1 - \alpha)c_{\theta_H}$  where  $\alpha$  is the fraction of the population that selects type  $\theta_L$ .

## 4.2 Equilibrium in the Infinite-Horizon Model

Equilibrium in the infinite-horizon model will also fall into one of the three cases illustrated informally for the static model: In each period, consumers will either all select the same type in equilibrium ( $\theta_L$  or  $\theta_H$ ), or there will be heterogeneity in the choice of type. We defer the details of how to solve for the pricing kernel and check the market-clearing conditions in each of these cases to Appendix A.2, but we present the numerical results of our equilibrium analysis in the following section.

## 5 Calibration

In this section, we numerically solve the infinite-horizon model. We will highlight several key features of individual risk attitudes and market equilibrium for various parameter values: heterogeneity in market participation and risk exposure, asset returns, and attitudes toward small and large idiosyncratic gambles.

### 5.1 Parameter Values

Recall that aggregate consumption growth is determined by a state which is i.i.d. across time. Assume there are two states,  $Z = \{z^l, z^h\}$ , and each occurs with equal probability:  $P(z^l) = P(z^h) = 0.5$  (we relax this assumption in Section 5.3). Aggregate consumption satisfies  $e_{t+1} = \lambda(z_{t+1})e_t$ . For ease of comparison to existing results, we calibrate the model using the same mean and standard deviation for aggregate consumption growth as in Mehra and Prescott (1985):

$$\lambda(z^l) = \mu - \sigma, \quad \lambda(z^h) = \mu + \sigma$$

for  $\mu = 1.018$  and  $\sigma = 0.036$ . We will also analyze the returns of an asset that pays a stream of dividends  $\{d_t\}$ . The growth rate of dividends satisfies  $d_{t+1} = \lambda^d(z_{t+1})d_t$ , where

$$\lambda^d(z^l) = \mu - \sigma_d, \quad \lambda^d(z^h) = \mu + \sigma_d$$

for  $\sigma_d = 0.10$ .

There is a continuum of consumers with identical preferences. Each consumer has a value function  $\mathcal{V}(w_t; M)$  for wealth (conditional on the pricing kernel) that satisfies Equations (1), (2), and (3). We consider the following simple specification:

$$\Theta = \{\theta_L, \theta_H\}, \quad \theta_H > \theta_L, \quad \tau(\theta) = \begin{cases} \tau_H = 0 & \text{if } \theta = \theta_H \\ \tau_L > 0 & \text{if } \theta = \theta_L. \end{cases}$$

Several values of  $\theta_L$ ,  $\theta_H$ , and  $\tau_L$  will be considered in Table 1.

As part of the calibration of the model, we also describe the attitudes toward 50-50 gambles of various scales for the different parameter values under consideration. Examining the gains needed to compensate for losses ranging from small to large provides another gauge of the overall risk attitudes associated with different specifications, and hence of whether parameter values generate reasonable behavior outside of this specific investment application.<sup>11</sup> Similar explorations have been used to evaluate the choice of parameters

---

<sup>11</sup>There are many estimates of what constitute reasonable values for the coefficient of relative risk

in other models (e.g., Epstein and Zin (1990); Kandel and Stambaugh (1991)).

Imagine a consumer is offered a one-time gamble over future wealth at some initial period  $t$ . Inserting the explicit formula for the value function from Equation (4) into Equation (3), the consumer evaluates the random future wealth  $w_{t+1}$  resulting from this gamble according to

$$\mathcal{R}(\mathcal{V}(w_{t+1}; M)) = \Lambda(M) + \max_{\theta \in \Theta} \left\{ \log \left( \mathbb{E}[w_{t+1}^{-\theta}]^{-\frac{1}{\theta}} \right) - \tau(\theta) \right\}. \quad (5)$$

We will use Equation (5) to evaluate atemporal wealth gambles in the calibration results that follow.<sup>12</sup>

## 5.2 Numerical Results

Table 1 summarizes our results. For comparison, the first two columns of the table describe EZKP utility with coefficients of relative risk aversion  $(\theta_H + 1)$  equal to 4 and 18, respectively. The last three columns describe different parameter values for optimal risk attitude utility.

Panel A in the table describes the gains needed for an individual to accept an atemporal 50-50 gamble for various possible loss values when initial wealth is \$300,000. Panel B describes the consumption growth rates of each type  $\theta \in \{\theta_L, \theta_H\}$  that is selected by some segment of the population in equilibrium. We write  $\lambda_\theta(z)$  to denote the consumption growth rate of type  $\theta$  consumers in the current period:  $\lambda_\theta(z) = c_{i,t+1}(z)/c_{i,t}$  for any consumer  $i$  who selects type  $\theta$  in period  $t$ . In addition to describing the mean and standard deviation of consumption growth for each type, the last row of this panel indicates the fraction of current aggregate wealth held by consumers who choose type  $\theta_L$  or, equivalently, the fraction of the current period aggregate endowment consumed by type  $\theta_L$  consumers.<sup>13</sup> Panel C indicates the equilibrium values of the pricing kernel and asset returns. In this panel,  $R^f$  denotes the gross risk-free rate and  $R$  denotes the gross return on an asset paying the stochastic dividend stream  $\{d_t\}$ .

Turning to the first specification in the table, note that EZKP1 generates reasonable aversion to the largest gambles in Panel A, but it is almost risk neutral for small gambles.

---

aversion in a CRRA expected-utility model; however, one cannot rely on these estimates since the risk attitude of individuals with ORA preferences may change with their exposure to risk.

<sup>12</sup>Epstein and Zin (1989, Section 5) were the first to demonstrate that, assuming the stochastic process driving the economy is i.i.d. across time, the certainty equivalent in a recursive non-expected-utility model also represents the preferences over timeless wealth gambles.

<sup>13</sup>Just as in the single-period illustration from Section 4.1, in a heterogeneous-type equilibrium in the infinite-horizon economy, the fraction  $\alpha$  of aggregate wealth held by consumers who choose type  $\theta_L$  is pinned down by the market clearing condition, in this case applied to consumption growth rates rather than levels:  $\lambda(z) = \alpha\lambda_{\theta_L}(z) + (1 - \alpha)\lambda_{\theta_H}(z)$  for  $z \in \{z^l, z^h\}$ . See Appendix A.2 for details.

**Table 1. Calibration Results: EZKP and ORA Models**

The table has results for two specifications of the EZKP model and three specifications of the ORA model. Consumption and dividend growth rates are summarized in Section 5.1,  $R^f$  denotes the risk-free rate, and  $R$  denotes the gross return on asset paying the stochastic dividend stream  $\{d_t\}$ .

Risk-Preference Model					
	EZKP1	EZKP2	ORA1	ORA2	ORA3
$\theta_H$	3.000	17.000	25.000	25.000	100.000
$\theta_L$	–	–	3.000	4.000	3.000
$\tau_L$	–	–	0.020	0.025	0.020
$\beta^{-1}$	1.010	1.010	1.010	1.010	1.010

Panel A: Binary 50-50 Gambles					
Loss	Gain that leads to indifference for initial wealth \$300,000				
\$100	100.13	100.60	100.87	100.87	103.48
\$400	402.14	409.84	414.37	414.37	462.37
\$1,000	1,013.51	1,063.85	1,094.95	1,094.95	1,518.31
\$5,000	5,357.20	7,170.61	8,995.81	8,995.81	18,991.43
\$10,000	11,539.60	27,901.22	26,396.79	32,281.88	26,396.79
\$20,000	27,302.60	$\infty$	45,692.48	58,228.29	45,692.48
\$30,000	50,274.57	$\infty$	75,052.03	110,405.61	75,052.03

Panel B: Equilibrium Consumption Growth by Type					
$\lambda_{\theta_H}(z^l)$	0.9820	0.9820	0.9927	0.9880	1.0011
$\lambda_{\theta_H}(z^h)$	1.0540	1.0540	1.0277	1.0342	1.0095
$\mathbb{E}(\lambda_{\theta_H})$	1.0180	1.0180	1.0102	1.0111	1.0053
$\sigma(\lambda_{\theta_H})$	0.0360	0.0360	0.0175	0.0231	0.0042
$\lambda_{\theta_L}(z^l)$	–	–	0.9346	0.9399	0.9374
$\lambda_{\theta_L}(z^h)$	–	–	1.1709	1.1919	1.1576
$\mathbb{E}(\lambda_{\theta_L})$	–	–	1.0527	1.0659	1.0475
$\sigma(\lambda_{\theta_L})$	–	–	0.1182	0.1260	0.1101
% type $\theta_L$	–	–	18.36	12.55	30.04

Panel C: Pricing Kernel and Asset Returns					
$M(z^l)$	1.1149	1.5508	1.4047	1.5192	1.3796
$M(z^h)$	0.8400	0.4339	0.5700	0.4633	0.5934
$R^f$	1.0231	1.0077	1.0128	1.0088	1.0137
$\mathbb{E}(R)$	1.0374	1.0667	1.0567	1.0645	1.0550
$\sigma(R)$	0.1019	0.1048	0.1038	0.1046	0.1036
$\mathbb{E}(R) - R^f$	0.0143	0.0590	0.0439	0.0557	0.0413



It also only generates an equity premium of 1.4%. In contrast, specification EZKP2 assumes a coefficient of relative risk aversion of 18 and generates a more realistic equity premium of 5.9%.<sup>14</sup> However, the gambling behavior for this specification highlights the concerns about assuming such a large coefficient of risk aversion. The individual will reject any gamble in which she will lose over \$20,000 (7% of wealth) with even odds, regardless of the size of the possible gain that could be won. These specifications of EZKP utility provide a parametric illustration of a paradoxical implication of expected utility that was shown by Rabin (2000) to hold more generally: Any expected-utility preference must have either implausibly low aversion to small gambles or excessive aversion to large gambles. Similar observations related to large-scale idiosyncratic risks (e.g., occupational earnings risk) led Mehra and Prescott (1985), Lucas (2003), and others to argue that the coefficient of relative risk aversion should be bounded above by 10.

Specification ORA1 improves the risk attitudes for binary gambles over both EZKP1 and EZKP2: It increases the risk aversion for small gambles relative to these specifications, while drastically reducing the extreme risk aversion of EZKP2 for large gambles. At the same time, it maintains a moderately large equity premium of 4.4%.<sup>15</sup> Panel B illustrates the heterogeneity in consumption risk for the two segments of the population under ORA1. Just over 18% of the aggregate wealth in the population is held by consumers who choose type  $\theta_L$ . The consumption growth rate  $\lambda_{\theta_L}$  for this fraction of the population has a standard deviation of 0.1182, over three times that of aggregate consumption growth. The majority of the population chooses type  $\theta_H$ , and the consumption growth rate  $\lambda_{\theta_H}$  for this segment of the population has a standard deviation of 0.0175, roughly half that of aggregate consumption. Thus a significant portion of aggregate risk is consolidated in a small segment of the population. Moreover, the additional risk borne by type  $\theta_L$  consumers is compensated by a substantially higher expected consumption growth rate: 1.05 (5%) rather than the expected consumption growth rate of 1.01 (1%) for type  $\theta_H$  consumers.<sup>16</sup>

These results are consistent with the observed patterns of investment summarized in the introduction. There has traditionally been a small segment of the population that invests in the stock market, with the majority of the population investing primarily in

---

<sup>14</sup>It is well known that EZKP utility with a high degree of risk aversion can generate a large equity premium. The calibration in specification EZKP2 is consistent with the estimate from Kocherlakota (1996, page 51) that a coefficient of relative risk aversion of 17.95 will satisfy the Euler equation for the equity premium.

<sup>15</sup>One caveat in interpreting these results is that the simple two-state stochastic process in our analysis implies perfect correlation between consumption and asset returns. Imposing more realistic correlation between aggregate consumption and returns tends to lower the equity premium in most models.

<sup>16</sup>Cross-sectional differences in expected returns due to differences in portfolio allocations have been suggested as one potential driver of wealth inequality. However, our model lends a very different welfare interpretation to these differences, since the portfolio choices of type  $\theta_H$  and  $\theta_L$  consumers yield the same ex-ante utility.

low-risk and low-return savings instruments. We should be careful to point out that our model predicts that even type  $\theta_H$  consumers continue to bear some consumption risk. However, as we discuss in Section 6.1, since the utility difference between the equilibrium allocation with minimal risk exposure and one with no exposure to stock market risk is relatively small, introducing modest participation costs could drive some of these consumers completely out of the market—even those with relatively high wealth.

The last two columns illustrate the impact of changing the parameters of ORA utility. Specification ORA2 increases the values of  $\theta_L$  and  $\tau_L$ . The result is an increase in the equity premium, up to 5.6%, and a decrease in the fraction of the population choosing type  $\theta_L$  and holding greater consumption risk, down to under 13%. For this specification, aversion to large gambles on the scale of \$30,000 (10% of wealth) becomes overly large, although still more reasonable than for EZKP2. Specification ORA3 illustrates the impact of instead increasing  $\theta_H$ . This generates a high level of risk aversion for small-scale gambles, but the attitude toward very large gambles is the same as under ORA1. It is interesting to note that due to changes in equilibrium consumption heterogeneity, the equity premium actually decreases slightly for this specification.


### 5.3 Comparative Statics of Participation and Risk Sharing

We now examine how risk sharing in the population responds to changes in fundamentals. There are many possible changes to aggregate uncertainty that have been considered in the literature, including changes to the conditional expectation and the conditional volatility of consumption growth. In this section, we continue to focus on a consumption growth process that is i.i.d. across time and examine a particular change to its (unconditional) distribution. To facilitate comparison to existing results, we maintain the same mean and standard deviation of consumption growth and only vary its skewness.

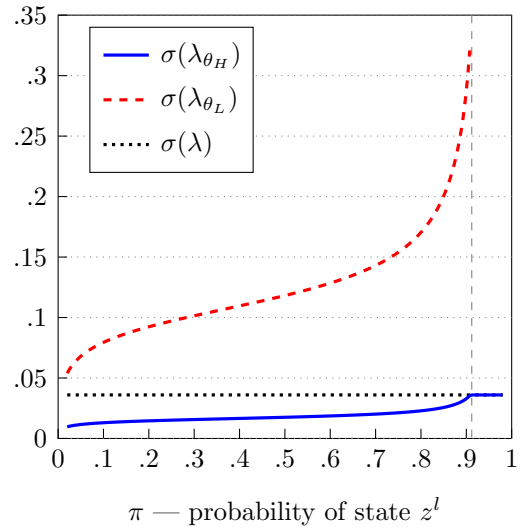
Let  $\pi$  denote the probability of state  $z^l$ . In the analysis in Table 1, this value was fixed at  $\pi = 0.5$ . We now consider the impact of varying this probability while holding fixed the first and second moments of consumption and dividend growth. Specifically, growth rates take the following values:

$$\begin{aligned} \lambda(z^l) &= \mu - \sigma \sqrt{\frac{1-\pi}{\pi}} & \lambda(z^h) &= \mu + \sigma \sqrt{\frac{\pi}{1-\pi}} \\ \lambda^d(z^l) &= \mu - \sigma_d \sqrt{\frac{1-\pi}{\pi}} & \lambda^d(z^h) &= \mu + \sigma_d \sqrt{\frac{\pi}{1-\pi}}. \end{aligned} \tag{6}$$

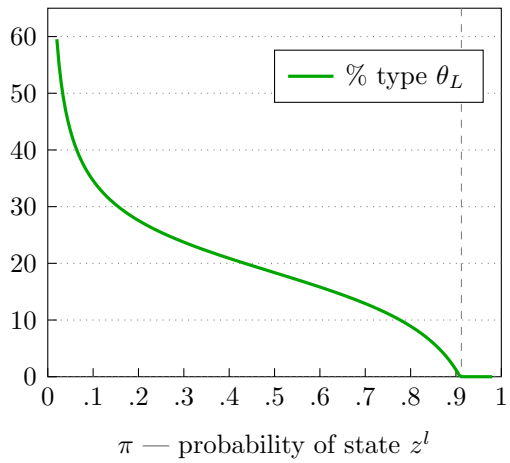
Using the parameter values from specification ORA1, Figure 3 illustrates the impact of changing skewness on the expected value and standard deviation of the consumption growth rates of each type, on the distribution of types in the population, and on asset




(a) Expected value of consumption growth rate by type



(b) Standard deviation of consumption growth rate by type



(c) Percent of aggregate wealth held by consumers choosing type  $\theta_L$  in a given period



(d) Expected returns and risk premium

**Figure 3.** Expected value and standard deviation of consumption growth by type, division of population into types, and returns. Each is plotted as a function of the probability  $\pi$  of state  $z^l$ . Consumption and dividend growth rates satisfy Equation (6), and parameters values are set according to the ORA1 specification in Table 1.

returns. As  $\pi$  increases, aggregate consumption growth transitions from being skewed to the left to being skewed to the right. In response to this change, the proportion of type  $\theta_L$  consumers (as measured by fraction of aggregate wealth held by this type) decreases while simultaneously the consumption risk held by each of these consumers increases. The standard deviation of  $\lambda_{\theta_L}$  spikes to roughly 9 times that of aggregate risk as  $\pi$  approaches 0.91, and once  $\pi$  exceeds this threshold all consumers select type  $\theta_H$  and hold identical portfolio allocations.

While these comparative statics concern changes in the unconditional distribution of an i.i.d. process, they are suggestive that the ORA model may provide a mechanism to understand intertemporal variation in investment levels and movements in and out of the stock market. If consumption growth follows a stationary Markov process with state-dependent mean, standard deviation, or skewness, one might expect the equilibrium participation rate to change with the current state. Extending our analysis to more general stochastic processes and relating its predictions to recent studies on the dynamics of household portfolio decisions is an obvious area for future exploration.

## 6 Discussion and Extensions

### 6.1 Participation Costs

As we noted earlier, the combination of optimal risk attitude utility with participation costs can be used to further refine the predictions of our model. The analysis in the previous section showed that consumers who select type  $\theta_H$  in equilibrium invest significantly less in stocks, and their utility gain from these investments are relatively small given their high level of local risk aversion (recall that utility for low-risk allocations is equivalent to that of an EZKP utility maximizer with a coefficient of relative risk aversion of  $\theta_H + 1$ ). This implies that the level of participation costs required to induce nonparticipation within the ORA model will be much lower than previous estimates based on expected-utility preferences.

An important consequence of this observation is that, even for wealthy households, much more moderate financial or informational costs can lead to nonparticipation, which allows our model to address one of the most puzzling facts from household finance: nonparticipation among the wealthy. Moreover, since the required costs are small, slight cross-sectional variation in risk aversion or participation costs can generate the combination of complete nonparticipation by some households, small but positive investment in stocks by some households, and high levels of investment by others, where the portion of wealth invested in stocks is positively (but not perfectly) correlated with wealth. We leave the formal analysis of this extension as an important direction for future research.

## 6.2 Idiosyncratic Income Risk

As we discussed in Section 2, one difficulty for models that rely on first-order risk aversion to explain nonparticipation is that the introduction of background risk tends to significantly decrease the induced aversion to stock market risk for such preferences, making it harder to rationalize nonparticipation using reasonable parameter values. This observation motivated Barberis, Huang, and Thaler (2006) to incorporate narrow framing of stock market risk into their model of participation.

Interestingly enough, this property of models with first-order risk aversion differs dramatically from the implications of background risk for expected utility: For CRRA expected-utility preferences, the introduction of independent background risk only serves to increase effective risk aversion and hence to decrease the optimal investment in stocks (see Gollier and Pratt (1996)). In fact, the asset-pricing literature has used incomplete consumption insurance (in the form of persistent idiosyncratic income shocks) to generate *greater* aversion to stock market risk in order to help explain the level of the equity premium (see Constantinides and Duffie (1996); Brav, Constantinides, and Geczy (2002)).

The same logic suggests that narrow framing is unlikely to be needed to generate realistic predictions with our model. Based on existing theoretical results for expected utility, it is easy to see that introducing small or moderate amounts of background risk into the ORA model will similarly increase aversion to stock market risk and decrease the optimal level of investment in stocks. Assuming the amount of background risk is not sufficiently large to change the optimal risk attitude  $\theta$ , consumers in our model behave exactly like CRRA expected-utility maximizers, and hence the comparative statics results in Gollier and Pratt (1996) apply. However, there is an important caveat to this argument: If the amount of background risk is sufficiently large, then consumers who would otherwise choose the risk attitude  $\theta_H$  may switch to choosing  $\theta_L$ , thereby decreasing their local risk aversion. Therefore, the impact of uninsurable idiosyncratic risk on our results will depend on the exact scale of the risk.

## 6.3 Optimal Expectations and Speculative Behavior

For ORA utility, the individual optimizes her risk attitude for any given distribution of future outcomes. There is a literature that considers the dual problem of optimizing beliefs for a fixed utility function (e.g., Brunnermeier and Parker (2005); Gollier and Muermann (2010); Bénabou and Tirole (2011); Macera (2014)). The models in this literature usually predict distortions of future behavior in response to changes in current beliefs, or require distortions of current behavior as a form of self-signaling. These features are in sharp contrast to our model of dynamically-consistent choice.

The differences between these two approaches allows each to address slightly different questions. For example, [Brunnermeier and Parker \(2005\)](#) showed that optimal expectations can generate endogenous heterogeneity in investment behavior in equilibrium, where ex-ante identical consumers select opposing beliefs and take stock market positions that bet against one another. Optimal expectations can thus help to explain speculative investment behavior. Our model addresses a different type of heterogeneity, namely, greater risk bearing by one segment of the population.

# A Additional Derivations

This section provides some supporting results that are used in the main text. First, we establish the existence and functional form of the value function for consumers' preferences. Second, we describe the precise equilibrium conditions for the economy for each of three possible cases and summarize our numerical procedure. Third, we provide the formulas for asset returns.

## A.1 Properties of the Value Function

Since equilibrium may involve heterogeneity in the choice of  $\theta$ , it is useful to define the value of the optimization problem conditional on choosing risk attitude  $\theta \in \Theta$  in a given period:

$$\mathcal{V}_\theta(w_t; M) = \max_{\substack{c_t \in \mathbb{R}_+ \\ w_{t+1} \in \mathbb{R}_+^Z}} \left\{ (1 - \beta) \log(c_t) - \beta \frac{1}{\theta} \log \left( \mathbb{E} \left[ \exp(-\theta \mathcal{V}(w_{t+1}; M)) \right] \right) - \tau(\theta) \right\}, \quad (7)$$

subject to the budget constraint in Equation (1). Thus

$$\mathcal{V}(w_t; M) = \max_{\theta \in \Theta} \mathcal{V}_\theta(w_t; M).$$

The following proposition summarizes the relevant properties of  $\mathcal{V}$  and  $\mathcal{V}_\theta$  that will be used in the equilibrium analysis in Section A.2. The first part of this result is precisely Proposition 1.

**Proposition 2** *For any pricing kernel  $M : Z \rightarrow \mathbb{R}_{++}$ , there exists a unique value function  $\mathcal{V}$  for the problem described in Equations (1), (2), and (3). This value function takes the form*

$$\mathcal{V}(w_t; M) = \Lambda(M) + \log(w_t),$$

where

$$\Lambda(M) = \log(1 - \beta) + \frac{\beta}{1 - \beta} \log(\beta) + \frac{\beta}{1 - \beta} \max_{\theta \in \Theta} \left\{ \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \tau(\theta) \right\}.$$

Moreover, the maximizing consumption and state-contingent future wealth in the conditional (on  $\theta \in \Theta$ ) optimization problem in Equation (7) are

$$\begin{aligned} c_{\theta,t} &= (1 - \beta)w_t, \\ w_{\theta,t+1}(z) &= \beta w_t \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-1} M(z)^{-\frac{1}{\theta+1}}, \end{aligned}$$

and therefore  $\mathcal{V}_\theta$  takes the form

$$\begin{aligned} \mathcal{V}_\theta(w_t; M) &= \beta \Lambda(M) + (1 - \beta) \log(1 - \beta) + \beta \log(\beta) \\ &\quad + \beta \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \beta \tau(\theta) + \log(w_t). \end{aligned}$$

The proof of Proposition 2 is contained in Appendix B. For this particular specification, the existence of the value function follows from the theorem of Blackwell (1965).<sup>17</sup>

## A.2 Equilibrium Conditions and Numerical Procedure

Equilibrium in the model falls into one of three cases. These correspond roughly to the three cases described in the static context in Section 4 and depicted in Figures 1 and 2, but adapted to deal with consumption growth rather than consumption levels. As we explain at the end of this section, checking each of these possible cases is precisely the numerical procedure used in the calibration in Section 5:

1. *All consumers choose  $\theta = \theta_L$  in equilibrium:* Combining the two optimality conditions from Proposition 2, the each consumer  $i \in [0, 1]$  must satisfy

$$c_{i,t+1}(z) = \beta c_{i,t} \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_L+1}}.$$

Aggregating across consumers, this implies the following relationship between aggregate consumption growth and the pricing kernel:

$$\lambda(z) = \beta \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_L+1}}.$$

The solution to this equation is<sup>18</sup>

$$M(z) = \beta \mathbb{E} [\lambda^{-\theta_L}]^{-1} (\lambda(z))^{-(\theta_L+1)}. \quad (8)$$

To check whether or not this is in fact an equilibrium, we need to determine whether  $\theta_L$  is actually optimal given these prices. Note that by Proposition 2,  $\mathcal{V}_{\theta_L}(w_t; M) \geq \mathcal{V}_{\theta_H}(w_t; M)$  holds if and only if

$$e^{-\tau_L} \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-\frac{\theta_L+1}{\theta_L}} \geq \mathbb{E} \left[ M^{\frac{\theta_H}{\theta_H+1}} \right]^{-\frac{\theta_H+1}{\theta_H}}. \quad (9)$$

If this condition is satisfied, then we have found an equilibrium. If not, then there is not an equilibrium in which all consumers choose  $\theta = \theta_L$ .

---

<sup>17</sup>It is worth noting that, as with most recursive non-expected-utility models, the usual techniques from Blackwell (1965) can only be applied in a very restricted set of special cases of the general ORA model from Sarver (2017). Fortunately, for other homogeneous specifications, recent results by Marinacci and Montrucchio (2010) can be used to establish the existence and uniqueness of the value function (see also Theorem 7 in Sarver (2012) which builds on their results).

<sup>18</sup>Equation (8) illustrates that when consumption growth is i.i.d. across time and all consumers select the same type  $\theta_L$ , the pricing kernel is the same as for time-separable expected utility with a coefficient of relative risk aversion  $\theta_L + 1$  and discount factor  $\hat{\beta} = \beta \mathbb{E} [\lambda^{-\theta_L}]^{-1}$ . This formula could also be obtained by appealing to the results of Kocherlakota (1990), who observed the equivalence of the equilibrium conditions for EZKP utility and time-separable expected utility (with a different discount factor) in i.i.d. environments.



2. *All consumers choose  $\theta = \theta_H$  in equilibrium:* The analysis is similar to the previous case. To satisfy the consumer optimality conditions given aggregate consumption growth, the pricing kernel must be

$$M(z) = \beta \mathbb{E}[\lambda^{-\theta_H}]^{-1} (\lambda(z))^{-(\theta_H+1)}. \quad (10)$$

If these prices satisfy  $\mathcal{V}_{\theta_H}(w_t; M) \geq \mathcal{V}_{\theta_L}(w_t; M)$ , which is equivalent to

$$\mathbb{E} \left[ M^{\frac{\theta_H}{\theta_H+1}} \right]^{-\frac{\theta_H+1}{\theta_H}} \geq e^{-\tau_L} \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-\frac{\theta_L+1}{\theta_L}}, \quad (11)$$

then we have found an equilibrium. If not, then there is not an equilibrium in which all consumers choose  $\theta = \theta_H$ .

3. *Consumers are heterogeneous, with some choosing  $\theta = \theta_L$  and some choosing  $\theta = \theta_H$  in equilibrium:* If both  $\theta_L$  and  $\theta_H$  are optimal, then they must give the same indirect utility. This gives our first equilibrium restriction on prices:  $\mathcal{V}_{\theta_L}(w_t; M) = \mathcal{V}_{\theta_H}(w_t; M)$  or, equivalently,

$$e^{-\tau_L} \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-\frac{\theta_L+1}{\theta_L}} = \mathbb{E} \left[ M^{\frac{\theta_H}{\theta_H+1}} \right]^{-\frac{\theta_H+1}{\theta_H}}. \quad (12)$$

We use individual optimality and market clearing conditions to give our second restriction. Let  $\alpha \in (0, 1)$  denote the fraction of time  $t$  wealth held by consumers who choose type  $\theta_L$ . Then the total time  $t$  consumption by type  $\theta_L$  consumers is  $\alpha e_t$ , and the total time  $t$  consumption by type  $\theta_H$  consumers is  $(1 - \alpha)e_t$ . By the consumers' optimality conditions, total consumption of these two groups at time  $t + 1$  are as follows:

$$\begin{aligned} \text{type } \theta_L \text{ consumers: } & \alpha \beta e_t \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_L+1}} \\ \text{type } \theta_H \text{ consumers: } & (1 - \alpha) \beta e_t \mathbb{E} \left[ M^{\frac{\theta_H}{\theta_H+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_H+1}}. \end{aligned}$$

Market clearing requires that these sum to  $e_{t+1}$  or, equivalently,

$$\lambda(z) = \alpha \beta \mathbb{E} \left[ M^{\frac{\theta_L}{\theta_L+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_L+1}} + (1 - \alpha) \beta \mathbb{E} \left[ M^{\frac{\theta_H}{\theta_H+1}} \right]^{-1} M(z)^{-\frac{1}{\theta_H+1}}. \quad (13)$$

If Equations (12) and (13) are both satisfied, then we have found an equilibrium.

Note that Equation (13) implies that  $\mathbb{E}[M\lambda] = \beta$ . In the two-state version of the model used in the calibration in Section 5 with  $Z = \{z^l, z^h\}$ , it can be shown that a partial converse is also true: If  $\mathbb{E}[M\lambda] = \beta$  then Equation (13) is satisfied for some  $\alpha \in \mathbb{R}$ . This observation permits the numerical problem for this case to be reduced to solving an equation of a single variable. Specifically, we adopt the following numerical procedure:

1. Check for an equilibrium as in case 1 as follows: Define  $M$  as in Equation (8) and check if Equation (9) is satisfied. If that fails, then proceed to step 2.

2. Check for an equilibrium as in case 2 as follows: Define  $M$  as in Equation (10) and check if Equation (11) is satisfied. If that fails, then proceed to step 3.
3. Find an equilibrium as in case 3 as follows: Since the solution to Equation (12) is only pinned down up to a constant, look for a solution of the form

$$\hat{M}(z) = \begin{cases} p & \text{if } z = z^l \\ 1 & \text{if } z = z^h, \end{cases}$$

for  $p \geq 1$ . Then let

$$M(z) = \frac{\beta}{\mathbb{E}[\hat{M}\lambda]} \hat{M}(z).$$

Since  $M$  is a scalar multiple of  $\hat{M}$  it satisfies Equation (12), and by construction it satisfies  $\mathbb{E}[M\lambda] = \beta$ . Thus it satisfies Equation (13) for some  $\alpha \in \mathbb{R}$ . Finally, it can be shown that if there is not an equilibrium as in case 1 or 2, then we necessarily have  $\alpha \in (0, 1)$  and hence the market clearing condition is satisfied. Thus we have found a heterogeneous-type equilibrium.

### A.3 Asset Returns

In this section, we derive the formulas for asset returns for the calibration in Section 5 using the equilibrium pricing kernel. By the definition, the period  $t$  price of any random asset paying  $x_{t+1}$  in period  $t + 1$  is  $p_t = \mathbb{E}_t[M_{t,t+1}x_{t+1}]$ . The risk-free rate is therefore given by

$$1 = \mathbb{E}[M_{t,t+1}R_t^f] \quad \text{or} \quad R_t^f = \frac{1}{\mathbb{E}_t[M_{t,t+1}]}.$$

Since the equilibrium pricing kernel is stationary and satisfies  $M_{t,t+1} = M(z_{t+1})$ , the risk-free rate is constant and is given by the unconditional expectation  $R^f = 1/\mathbb{E}[M]$ .

Consider now an asset that has a dividend growth given by  $d_{t+1} = \lambda^d(z_{t+1})d_t$ . The price of this asset satisfies

$$p_t^d = \mathbb{E}_t[M_{t,t+1}(p_{t+1}^d + d_{t+1})].$$

The price/dividend ratio therefore satisfies

$$\frac{p_t^d}{d_t} = \mathbb{E}_t \left[ M_{t,t+1} \left( \frac{p_{t+1}^d}{d_{t+1}} \frac{d_{t+1}}{d_t} + \frac{d_{t+1}}{d_t} \right) \right].$$

Since the state is i.i.d. across time and the equilibrium pricing kernel is given by  $M_{t,t+1} = M(z_{t+1})$  at every period  $t$ , there is a constant price/dividend ratio  $p$  that solves this equation:

$$p = \mathbb{E}[M\lambda^d(p + 1)] \implies \frac{p}{p + 1} = \mathbb{E}[M\lambda^d].$$

The return on this asset is given by

$$R_{t+1} = \frac{p_{t+1}^d + d_{t+1}}{p_t^d} = \frac{\frac{p_{t+1}^d}{d_{t+1}} \frac{d_{t+1}}{d_t} + \frac{d_{t+1}}{d_t}}{\frac{p_t^d}{d_t}}.$$

The return is therefore time invariant and satisfies

$$R(z) = \lambda^d(z) \left( \frac{p+1}{p} \right) = \frac{\lambda^d(z)}{\mathbb{E}[M\lambda^d]}.$$

## B Proof of Proposition 2

Note that for any function  $f : \mathbb{R}_+^Z \rightarrow \mathbb{R}$  and any  $\alpha \in \mathbb{R}$ , the risk-adjustment operator in Equation (3) satisfies  $\mathcal{R}(f + \alpha) = \mathcal{R}(f) + \alpha$ . Therefore, the value function operator defined by the right side of Equation (2) satisfies the conditions of the theorem of Blackwell (1965), which ensures the existence and uniqueness of the value function  $\mathcal{V}(w_t; M)$  for any  $M : Z \rightarrow \mathbb{R}_{++}$ .

We will verify that the functional form  $\mathcal{V}(w_t; M) = \Lambda(M) + \log(w_t)$  satisfies the recursion in Equation (2). Assuming this form of the value function and substituting into Equation (7) gives

$$\mathcal{V}_\theta(w_t) = \max_{\substack{c_t \in \mathbb{R}_+ \\ w_{t+1} \in \mathbb{R}_+^Z}} \left\{ (1 - \beta) \log(c_t) + \beta \log \left( \mathbb{E}[w_{t+1}^{-\theta}]^{-\frac{1}{\theta}} \right) - \beta\tau(\theta) + \beta\Lambda(M) \right\}, \quad (14)$$

where the maximization is subject to Equation (1). Applying standard Lagrangian optimization techniques to this problem gives the following maximizers:

$$\begin{aligned} c_{\theta,t} &= (1 - \beta)w_t, \\ w_{\theta,t+1}(z) &= \beta w_t \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-1} M(z)^{-\frac{1}{\theta+1}}. \end{aligned}$$

Note that these solutions imply

$$\begin{aligned} \mathbb{E}[w_{\theta,t+1}^{-\theta}]^{-\frac{1}{\theta}} &= \beta w_t \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-1} \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{1}{\theta}} \\ &= \beta w_t \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}}. \end{aligned}$$

Substituting these solutions into Equation (14) therefore yields

$$\begin{aligned} \mathcal{V}_\theta(w_t; M) &= (1 - \beta) \log((1 - \beta)w_t) + \beta \log(\beta w_t) \\ &\quad + \beta \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \beta\tau(\theta) + \beta\Lambda(M) \\ &= \beta\Lambda(M) + (1 - \beta) \log(1 - \beta) + \beta \log(\beta) \\ &\quad + \beta \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \beta\tau(\theta) + \log(w_t). \end{aligned}$$

This will establish that  $\mathcal{V}_\theta$  has the form claimed in the statement of the proposition, once it is verified that  $\mathcal{V}$  takes the form assumed above. To see that the latter is true, note that

$$\begin{aligned}\mathcal{V}(w_t; M) &= \max_{\theta \in \Theta} \mathcal{V}_\theta(w_t; M) \\ &= \beta \Lambda(M) + (1 - \beta) \log(1 - \beta) + \beta \log(\beta) \\ &\quad + \beta \max_{\theta \in \Theta} \left\{ \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \tau(\theta) \right\} + \log(w_t).\end{aligned}$$

Therefore, the assumed form  $\mathcal{V}(w_t; M) = \Lambda(M) + \log(w_t)$  satisfies the recursion in Equation (2) if

$$\Lambda(M) = \log(1 - \beta) + \frac{\beta}{1 - \beta} \log(\beta) + \frac{\beta}{1 - \beta} \max_{\theta \in \Theta} \left\{ \log \left( \mathbb{E} \left[ M^{\frac{\theta}{\theta+1}} \right]^{-\frac{\theta+1}{\theta}} \right) - \tau(\theta) \right\}.$$

We have thus established all of the claims in the statement of the proposition.

## References

- Ang, A., G. Bekaert, and J. Liu (2005): “Why Stocks May Disappoint,” *Journal of Financial Economics*, 76, 471–508. [7]
- Attanasio, O. P., J. Banks, and S. Tanner (2002): “Asset Holding and Consumption Volatility,” *Journal of Political Economy*, 110, 771–792, [5]
- Aumann, R. J. (1966): “Existence of Competitive Equilibria in Markets with a Continuum of Traders,” *Econometrica*, 34, 1–17. [12]
- Barberis, N., M. Huang, and R. H. Thaler (2006): “Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing,” *American Economic Review*, 96, 1069–1090. [7,21]
- Barillas, F., L. P. Hansen, and T. J. Sargent (2009): “Doubts or Variability?” *Journal of Economic Theory*, 144, 2388–2418. [3]
- Bénabou, R., and J. Tirole (2011): “Identity, Morals and Taboos: Beliefs as Assets,” *Quarterly Journal of Economics*, 126, 805–855. [21]
- Blackwell, D. (1965): “Discounted Dynamic Programming,” *Annals of Mathematical Statistics*, 36, 226–235. [24,27]
- Boyle, P., L. Garlappi, R. Uppal, and T. Wang (2012): “Keynes Meets Markowitz: The Trade-Off Between Familiarity and Diversification,” *Management Science*, 58, 253–272. [7]
- Brav, A., G. M. Constantinides, and C. C. Geczy (2002): “Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence,” *Journal of Political Economy*, 110, 793–824. [5,21]
- Briggs, J. S., D. Cesarini, E. Lingqvist, and R. Östling (2015): “Windfall Gains and Stock Market Participation,” NBER working paper. [6]
- Brunnermeier, M., and J. Parker (2005): “Optimal Expectations,” *American Economic Review*, 95, 1092–1118. [21,22]
- Campbell, J. Y. (2003): “Consumption-Based Asset Pricing,” in G. M. Constantinides, M. Harris, and R. Stulz (Eds.), *Handbook of the Economics of Finance*, Vol. 1, Elsevier. [5]
- Campbell, J. Y. (2006): “Household Finance,” *Journal of Finance*, 61, 1553–1604. [2]
- Cao, H. H., T. Wang, and H. H. Zhang (2005): “Model Uncertainty, Limited Market Participation, and Asset Prices,” *Review of Financial Studies*, 18, 1219–1251. [8]
- Chapman, D. A., and V. Polkovnichenko (2009): “First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes,” *Journal of Finance*, 64, 1863–1887. [8]
- Constantinides, G. M., and D. Duffie (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104, 219–240. [21]
- Dow, J., and S. R. Werlang (1992): “Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio,” *Econometrica*, 60, 197–204. [7]
- Epstein, L. G., and J. Miao (2003): “A Two-Person Dynamic Equilibrium Under Ambiguity,” *Journal of Economic Dynamics and Control*, 27, 1253–1288. [7]

- Epstein, L. G., and M. Schneider (2007): “Learning Under Ambiguity,” *Review of Economic Studies*, 74, 1275–1303. [7]
- Epstein, L. G., and T. Wang (1994): “Intertemporal Asset Pricing Under Knightian Uncertainty,” *Econometrica*, 62, 283–322. [7]
- Epstein, L. G., and S. Zin (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969. [3,10,15]
- Epstein, L. G., and S. Zin (1990): “‘First-Order’ Risk Aversion and the Equity Premium Puzzle,” *Journal of Monetary Economics*, 26, 387–407. [15]
- Epstein, L. G., and S. Zin (1991): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *Journal of Political Economy*, 99, 263–286. [3]
- Gollier, C., and A. Muermann (2010): “Optimal Choice and Beliefs with Ex Ante Savoring and Ex Post Disappointment,” *Management Science*, 56, 1272–1284. [21]
- Gollier, C., and J. W. Pratt (1996): “Risk Vulnerability and the Tempering Effect of Background Risk,” *Econometrica*, 64, 1109–1123. [21]
- Gomes, F., and A. Michaelides (2005): “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *Journal of Finance*, 60, 869–904. [6]
- Guiso, L., and P. Sodini (2013): “Household Finance: An Emerging Field,” in G. M. Constantinides, M. Harris and R. M. Stulz (Eds.), *Handbook of the Economics of Finance*, Volume 2, Amsterdam, The Netherlands: Elsevier. [2,7]
- Gul, F. (1991): “A Theory of Disappointment Aversion,” *Econometrica*, 59, 667–686. [7]
- Haliassos, M., and C. C. Bertaut (1995): “Why Do So Few Hold Stocks?,” *The Economic Journal*, 105, 1110–1129. [2]
- Hansen, L. P., and T. J. Sargent (2001): “Robust Control and Model Uncertainty,” *American Economic Review*, 91, 60–66. [3]
- Hansen, L. P., T. J. Sargent, and T. Tallarini (1999): “Robust Permanent Income and Pricing,” *Review of Economic Studies*, 66, 873–907. [3]
- Heaton, J., and D. Lucas (2000): “Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk,” *Journal of Finance*, 55, 1163–1198. [2]
- Kandel, S., and R. F. Stambaugh (1991): “Asset Returns and Intertemporal Preferences,” *Journal of Monetary Economics*, 27, 39–71. [15]
- Kocherlakota, N. (1990): “Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result,” *Journal of Finance*, 45, 175–190. [24]
- Kocherlakota, N. (1996): “The Equity Premium: It’s Still a Puzzle,” *Journal of Economic Literature*, 34, 42–71. [2,3,17]
- Kreps, D., and E. Porteus (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46, 185–200. [3,10]
- Lucas, R. E. (2003): “Macroeconomic Priorities,” *American Economic Review*, 93, 1–14. [3,5,17]

- Macera, R. (2014): “Dynamic Beliefs,” working paper. [21]
- Maenhout, P. J. (2004): “Robust Portfolio Rules and Asset Pricing,” *Review of Financial Studies*, 17, 951–983. [3]
- Mankiw, N. G., and S. P. Zeldes (1991): “The Consumption of Stockholders and Nonstockholders,” *Journal of Financial Economics*, 29, 97–112. [2,5]
- Marinacci, M., and L. Montrucchio (2010): “Unique Solutions for Stochastic Recursive Utilities,” *Journal of Economic Theory*, 145, 1776–1804. [24]
- Mehra, R., and E. Prescott (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15, 145–161. [2,3,14,17]
- Rabin, M. (2000): “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 68, 1281–1292. [3,17]
- Safra, Z., and U. Segal (2008): “Calibration Results for Non-Expected Utility Theories,” *Econometrica*, 76, 1143–1166. [7]
- Sarver, T. (2012): “Optimal Reference Points and Anticipation,” CMS-EMS Discussion Paper 1566, Northwestern University. [24]
- Sarver, T. (2017): “Mixture-Averse Preferences,” working paper. [2,4,9,10,11,24]
- Segal, U., and A. Spivak (1990): “First Order versus Second Order Risk Aversion,” *Journal of Economic Theory*, 51, 111–125. [6]
- Skiadas, C. (2003): “Robust Control and Recursive Utility,” *Finance and Stochastics*, 7, 475–489. [3]
- Starr, R. M. (1969): “Quasi-Equilibria in Markets with Non-Convex Preferences,” *Econometrica*, 37, 25–38. [12]
- Strzalecki, T. (2011): “Axiomatic Foundations of Multiplier Preferences,” *Econometrica*, 79, 47–73. [3]
- Tallarini, T. (2000): “Risk-Sensitive Real Business Cycles,” *Journal of Monetary Economics*, 45, 507–532. [3]
- Vissing-Jørgensen, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110, 825–853. [5]
- Vissing-Jørgensen, A. (2003): “Perspectives on Behavioral Finance: Does ‘Irrationality’ Disappear with Wealth? Evidence from Expectations and Actions,” *NBER Macroeconomics Annual*, 139–208. [6]
- Vissing-Jørgensen, A., and O. P. Attanasio (2003): “Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion,” *American Economic Review*, 93, 383–391. [5]
- Weil, P. (1989): “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, 24, 401–421. [3]
- Weil, P. (1990): “Nonexpected Utility in Macroeconomics,” *Quarterly Journal of Economics*, 105, 29–42. [3]