THE JOURNAL OF FINANCE • VOL. LXIV, NO. 4 • AUGUST 2009

First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes

DAVID A. CHAPMAN and VALERY POLKOVNICHENKO*

ABSTRACT

We examine a wide range of two-date economies populated by heterogeneous agents with the most common forms of nonexpected utility preferences used in finance and macroeconomics. We demonstrate that the risk premium and the risk-free rate in these models are sensitive to ignoring heterogeneity. This follows because of endogenous withdrawal by nonexpected utility agents from the market for the risky asset. This finding is important precisely because these alternative preferences have frequently been proposed as possible resolutions to various asset pricing puzzles, *and* they have all been examined exclusively in a representative agent framework.

The well-documented shortcomings of asset pricing models based, in part, on the expected utility (EU) hypothesis has motivated a large and rapidly growing set of pricing models that incorporate nonexpected utility preferences. The pricing implications of these models in the existing literature have been based exclusively on the assumption of a single representative agent (RA) who prices assets using the aggregate endowment as the consumption process. While aggregation is well understood in the standard EU setting (see Constantinides (1982)), it can be more complicated when some (or all) agents do not have EU. We provide a simple and intuitive examination of the asset pricing implications of heterogeneity for the most commonly used alternatives to EU.

Our primary finding is that the risk premium and the risk-free rate are sensitive to ignoring heterogeneity in all of the simple two-date, non-EU models that we examine. We compare the prices in a model with two-agent types to the prices in a model with a single representative agent who has a wealth-weighted average of the utility parameters in the two-agent economy. The economies are otherwise identical; for example, they have the same endowment process, market structure, etc. It is common for the single-agent model to overstate the equity premium, often by as much as 100% or more of the value in the two-agent model, and to understate the level of the risk-free rate, often by as much as 20% of the value in the two-agent model.

*Chapman is at Boston College and Polkovnichenko is at the University of Texas at Dallas. Polkovnichenko acknowledges financial support from a Business and Economics Research Grant of the McKnight Foundation at the University of Minnesota. We would like to thank an anonymous referee; Nick Barberis; John Campbell; Wayne Ferson; Iraj Fooladi; Evan Gatev; Campbell Harvey (the editor); Erzo Luttmer, Maureen O'Hara; Neil Pearson; Bryan Routledge; Jeremy Stein; and workshop participants at Boston College, Harvard University, McGill University, and the 2005 WFA meetings for helpful comments.

The preferences that we examine include all of the main alternatives to EU that have been considered in the asset pricing literature, and they are united by the property of *first-order risk aversion*. This feature of utility functions, as discussed in detail by Segal and Spivak (1990), implies that risk premiums are proportional to the standard deviation of returns, as opposed to the EU case where risk premiums are proportional to the variance of returns. In particular, we consider rank-dependent expected utility (RDEU), disappointment aversion (DA) and generalized disappointment aversion (GDA), ambiguity aversion (AA), and a general form of reference-dependent utility. When agents have any of these preferences, there are endogenous changes in their willingness to participate in the market for the risky asset. As one agent holds all of the risky securities in certain states of the world, the equivalence between the heterogenous agent economies and the weighted-average agent economies breaks down.

The conclusion that we draw from this broad set of examples is that the results from a representative agent analysis can be misleading about even basic asset pricing implications in economies with non-EU preferences. The endogenous participation decisions—and full insurance equilibria—are qualitatively and quantitatively important for asset pricing, and they need to be explicitly incorporated into these, increasingly common, non-EU models.

The paper is organized as follows. In Section I, we review the set of preference specification used in our analysis. The asset pricing framework is described in Section II, and Section III presents the model calibration. Our main results are in Section IV, with a separate analysis of reference-dependent preferences in Section V. Section VI contains our conclusions.

I. First-Order Risk Averse Preferences

Segal and Spivak (1990) define a preference relation as exhibiting first-order risk aversion if the implied risk premium on a small gamble does not vanish as the size of the gamble converges to zero.[1] These preferences have indifference curves that are not differentiable at some point(s) in the choice space. Segal and Spivak (1990) show that, from an empirical perspective, a first-order risk averse investor has two appealing properties relative to a second-order risk averse investor. Specifically, a first-order risk averse investor will purchase full insurance, whereas a second-order risk averse investor will not, and a first-order risk averse investor may not invest in a risky asset with a positive risk premium if the risk premium is small enough, while a second-order risk averse investor always allocates at least some wealth to this risky asset. This second result will have far-reaching implications for the models that we examine.

[1] Formally, if $\pi(t)$ is the amount of money that an agent would pay to avoid the gamble $x + t\tilde{\varepsilon}$, where $E[\tilde{\varepsilon}] = 0$, then $\pi(t)$ is $O(t)$ and $o(t)$ for first and second-order risk averse preferences, respectively.

In defining the versions of different utility functions that we will use throughout this paper, we make a number of simplifying assumptions in order to focus on the economic intuition. There are N possible states of nature labeled as $\{\omega_i\}_{i=1}^N$. The choice set is given by consumption lotteries, which are defined as consumption-probability pairs, $\{\mathbf{c}, \mathbf{p}\} \in \mathbb{R}_{++}^N \times \Delta_+^N$,[2] where Δ_+^N denotes the unit simplex in the non-negative orthant in $\mathbb{R}^N$. The first element of the pair corresponds to the level of consumption received in each state and the second element to the probability distribution over the states. Without loss of generality, we assume that the states are ordered from the lowest consumption level to the highest. $\mathbf{P}$ defines a vector of cumulative probabilities corresponding to $\mathbf{p} = [p_1, \ldots, p_N]'$. In the following subsections, we provide a brief description of the forms of non-EU preferences that we will examine.

A. Rank-Dependent Expected Utility

Rank-dependent models construct decision weights that are nonlinear functions of the cumulative distribution of all possible outcomes, rather than assigning objective probability weights to each possible state of the world. In order to construct these weights, the consumer must be able to rank all outcomes from the best to the worst. The most common specifications of such utilities are the *rank-dependent expected utility* (or *anticipated utility*) model of Quiggin (1982) and Yaari's (1987) *dual theory of choice*. Epstein and Zin (1990) is an early attempt to use first-order risk aversion to resolve the equity premium puzzle. They embed an RDEU specification in a recursive intertemporal utility functional. In a representative agent setting, the aggregate consumption process and RDEU preferences can generate an equity premium that is approximately one-third of the historical average. Therefore, RDEU represents a partial resolution of the equity premium puzzle.[3] Our presentation of the basic RDEU setup follows Quiggin (1993).

The RDEU value function assigned to a lottery $\{\mathbf{c}, \mathbf{p}\}$ is

$$U(\{\mathbf{c}; \mathbf{P}\}) = \sum_{k=1}^{N} q_k(\mathbf{P}) u(c_k), \tag{1}$$

where $u(\cdot) : \mathbb{R}_{++}^N \to \mathbb{R}$ is a strictly increasing function that represents the consumer's preferences over deterministic consumption outcomes, and $q_k(\cdot) > 0$ are decision weights. The weights are constructed using a strictly increasing and differentiable function $Z(\cdot) : [0, 1] \to [0, 1]$, where $Z(0) = 0$ and $Z(1) = 1$ are defined on the cumulative probability of the events P_i as follows:

$$q_i = Z(P_i) - Z(P_{i-1}), \tag{2}$$

for $i = 2, \ldots, N$, $P_0 = 0$, and $P_N = 1$.

[2] We will use **bold** typeface to denote vectors (lower case) or matrices (upper case).

[3] Polkovnichenko (2005) applies RDEU to examine the portfolio choices of an individual investor. He demonstrates that for a reasonably parameterized model an RDEU investor will choose portfolio allocations that are consistent with data on individual portfolio holdings.

The decision weight in each state depends on both the state's ranking and the entire probability distribution of possible outcomes. The precise form of $Z(\cdot)$ determines how the consumer transforms objective probabilities into subjective decision weights, and we will use the term *decision weight function* when referring to it. The weighting function disentangles the consumer's risk attitude from his marginal utility of wealth, which is captured in u.[4] The overweighting of "bad" events relative to their objective probability (pessimism or risk aversion) is implied by $Z(P) > P$, and the overweighting of "good" events (optimism or risk seeking) is implied by the reverse inequality. In particular, a concave Z corresponds to a globally pessimistic agent and a convex Z to a globally optimistic one.

We assume that the weighting function has the concave form

$$Z(P_i) = P_i^{\phi}, \tag{3}$$

for $0 < \phi \leq 1$. A concave Z implies positive risk premiums, which is economically plausible.[5] Equation (3) has the following additional advantages for our analysis: (i) if $\phi = 1$, the decision weights reduce to the actual probabilities and RDEU collapses to expected utility; (ii) it overweights extreme bad outcomes, which represents a form of risk aversion; and (iii) it uses only one parameter. Epstein and Zin (1990) consider this form of the weighting function in their single-agent economy. In our quantitative experiments, we will consider economies with heterogeneous consumers who have different degrees of pessimism, ϕ.

B. Disappointment Aversion

Disappointment averse (DA) preferences have been formulated by Gul (1991) and generalized by Routledge and Zin (2004). The economic intuition behind DA preferences is that they overweight outcomes that are "disappointing," which are defined as being below the certainty equivalent. Epstein and Zin (2001) estimate a representative agent with DA preferences in a dynamic utility functional using generalized method of moments and standard aggregate consumption data. They argue that the fitted DA model produces more intuitively plausible attitudes toward risk when compared to the best-fit EU model, and the DA model is not rejected by conventional measures of overall model fit. Bekaert, Hodrick, and Marshall (1997), however, find that DA preferences cannot fully explain the predictability of excess returns when they are embedded in a monetary two-country general equilibrium model.

Routledge and Zin (2004) generalize this model by allowing for a more general definition of the disappointment threshold, which may be above or below the certainty equivalent. They show that a calibrated general equilibrium model

[4] Yaari's theory is a special case that assumes $u(c) = c$. So, marginal utility of wealth is constant, but it admits risk averse behavior via concave decision weights.

[5] In addition, concavity of the weighting function, along with other assumptions made explicit below, is sufficient to prove the existence of a representative agent in a setting with heterogeneous preferences.

with these preferences can generate countercyclical risk aversion, although the calibrated model still does not resolve the Mehra-Prescott (1985) puzzle. Let $\{\mathbf{c}, \mathbf{p}\} \in \mathbb{R}^N_{++} \times \Delta^N_+$ be a consumption lottery, as defined earlier. We follow Routledge and Zin (2004) in defining generalized disappointment aversion (GDA) and the certainty equivalent of $\{\mathbf{c}, \mathbf{p}\}$ as the pair $\{u(\cdot), \mu(\mathbf{p})\}$ that solves

$$u(\mu(\mathbf{p})) = \sum_{n=1}^{N} p_n u(c_n) - \theta \sum_{c_n \leq \delta\mu(\mathbf{p})} p_n \{u[\delta\mu(\mathbf{p})] - u(c_n)\}, \tag{4}$$

where $u(\cdot)$ is the utility over outcomes function, as in the RDEU specification, and $\mu(\mathbf{p})$ is the certainty equivalent, δ and θ are preference parameters that define generalized disappointment aversion, $\delta\mu(\mathbf{p})$ defines the "disappointment threshold," and θ defines the utility penalty imposed by disappointing outcomes. The innovation in Routledge and Zin (2004) is to allow for the disappointment to be different from the certainty equivalent. If $\delta = 1$, then the GDA collapses to the simple DA specified in Gul (1991). If $\theta = 0$, then the GDA preferences collapse to EU. The risk aversion of the GDA utility function depends on both δ and θ.

C. Ambiguity Aversion

Ambiguity refers to uncertainty over the possible distribution of outcomes. This is distinct from *risk*, as in our previous examples, where agents are assumed to know the exact subjective distribution of states at date 1. Ellsberg (1961) argues that ambiguity affects choice and, in particular, typical behavior points to ambiguity aversion. Gilboa and Schmeidler (1989) formalize ambiguity aversion preferences in an atemporal setting, and Epstein and Schneider (2006) extend their framework to a dynamic environment. Epstein and Schneider (2006) develop a representative agent model with a dynamic version of AA preferences. They find that, qualitatively, ambiguity demands compensation and generates features of equilibrium returns that are consistent with some recent empirical findings. Specifically, their model can generate expected excess returns that increase in idiosyncratic volatility, and they can also generate skewness in measured excess returns.[6]

Gilboa and Schmeidler (1989) preferences, also known as Knightian uncertainty or max–min preferences, assume that agents have a set of prior beliefs

[6] AA preferences have also been used to examine portfolio choice in Dow and de Costa Werlang (1992). They show that, in a one-period setting, the portfolio choices of an investor with AA preferences will be characterized in part by a region where the investor chooses to hold none of the risky asset (either long or short). The size of this zero-holding region depends only on the investor's attitudes toward uncertainty and on his subjective prior distribution of the likelihood of future states, and not on the curvature of the utility function. Ang, Bekaert, and Liu (2005) extend the nonparticipation results of Dow and de Costa Werlang (1992) to DA preferences. They also show that a dynamic version of DA preferences can generate reasonable portfolio allocations with horizon effects that are difficult to reproduce in an EU model that relies solely on curvature of the utility over outcomes.

about possible distributions of outcomes $\{p_s\} \in \Pi$. The utility function is defined using the *worst* possible probability distribution from among those in Π:

$$U(\mathbf{c}; \mathbf{p}) = \min_{p_s \in \Pi} \sum_s p_s u(c_s).$$

When Π is a nonsingleton subset of the unit simplex, the agent chooses the distribution that places the highest weights on the worst outcomes. At the allocations where any two state-outcomes are equal, the indifference curves may exhibit a kink if within the set Π there are distributions potentially placing different weights on the two equal states. In particular, with two states the kinks are at the 45-degree line. The first-order risk aversion implies that agents may abstain from taking positions betting on ambiguous states of nature. Since individual agents may have different sets of priors, Π, in equilibrium, the agents with more pessimistic views of possible distributions may avoid participating in markets for risky assets.

D. Summary

As the previous discussion indicates, the different forms of first-order risk averse preferences we consider have different motivations in the psychology and decision-theory literatures, and there is no single encompassing functional form for utility that nests all of the models as parameter restrictions of the general structure. Versions of all of these utility functions require modification of the independence axiom and are therefore potentially consistent with some of the experimental evidence against EU. There are also differing levels of tractability for model estimation. For example, Routledge and Zin (2004) emphasize that GDA preferences retain linearity in the probabilities, which leads to first-order conditions for portfolio choice that are similar to standard Euler equations from time-separable EU models.

Nonetheless, these preferences do have important points in common. They all emphasize the important—and distinct—roles of utility over outcomes and attitudes about the probability of possible states in the future. The indifference curves for these different preference functions are not differentiable at all points in the choice space.

Figure 1 provides a graphical depiction of the indifference curves for DA and GDA utility in a simple two-state example; see also Figure 2 in Routledge and Zin (2004). In this two-state setting (for equivalent parameterizations) the indifference curves for RDEU and AA are the same as the DA curve. The axes show different levels of consumption in the two states and the 45-degree line defines equal consumption across states. In the DA (RDEU and AA) cases, there is a single kink at the line of equal state consumption. The GDA indifference curves allow for multiple kinks representing a range of disappointing outcomes. It is important to note that the graphical equivalence between RDEU, AA, and DA does not extend to the multi-state setting considered later. When there are more than two states, the indifference curves for DA have kinks only at allocations

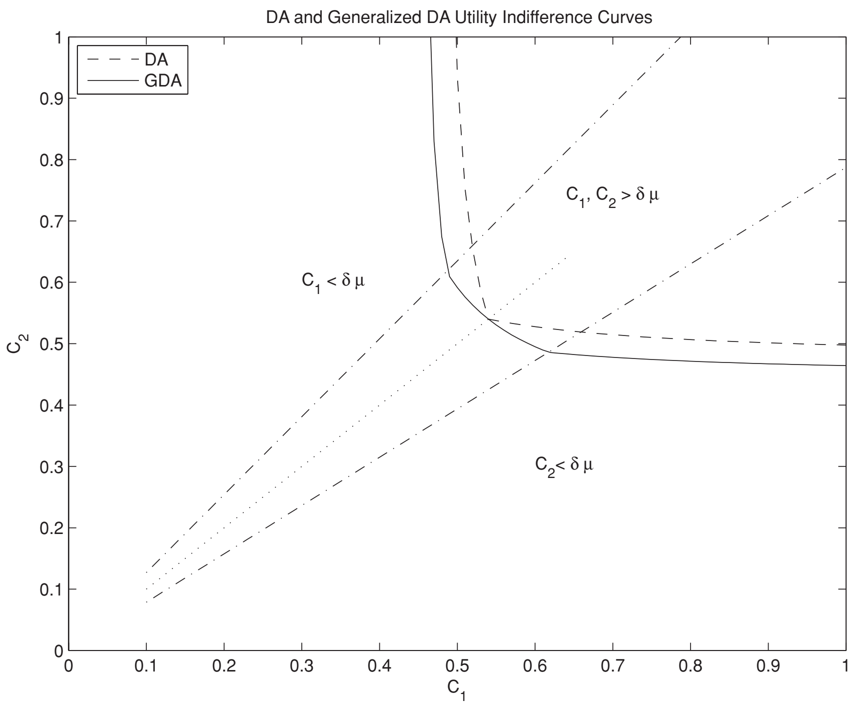

Figure 1. Indifference curves for a two-state version of DA and GDA utility functions.

where consumption is equal in all states (riskless), while the indifference curves for RDEU and AA have kinks at all allocations where consumption is equal in any two states (but not necessarily riskless across all states).

II. The Asset Pricing Framework

We use a common market structure where the only variation in our analysis is the specification of utility. The setting is deliberately simple. There are two dates, 0 and 1. The state at date 0 is known, and there are N possible states at date 1, labeled $\{\omega_i\}_{i=1}^{N}$.[7] These states are ordered from the lowest consumption level to the highest. The endowment at date 0 is known and denoted e_0. The endowment in each date 1 state is denoted $e(\omega_i)$, or e_i for short. The vector $\mathbf{e}_1 \in \mathbb{R}_+^N$ is the entire date 1 endowment vector. We assume that markets are complete and that there are no arbitrage opportunities. The date 0 prices of N contingent claims are denoted $\boldsymbol{\pi} \in \mathbb{R}_{++}^N$.

[7] The results in the following sections are not dependent on the simple assumptions that we make regarding the market structure. It is certainly possible to consider an infinite horizon model and more than two agents, and the two-period utility functions that we use here can be generalized to a dynamic structure following the approach taken in the various papers cited in the previous section.

An individual agent's asset pricing problem is standard

$$\max_{c_0,\boldsymbol{\theta}}\{u(c_0)+U(\mathbf{c_1};\mathbf{p})\} \tag{5}$$

subject to the budget constraints

$$c_0+\boldsymbol{\theta}'\boldsymbol{\pi}=e_0 \qquad \text{and} \qquad \mathbf{c}_1=\mathbf{e}_1+\mathbf{X}\boldsymbol{\theta}, \quad c_0\geq 0, c_1\geq 0, \tag{6}$$

where $U(\mathbf{c};\mathbf{p})$ represents one of the first-order risk averse utility functions defined earlier, $\boldsymbol{\theta}$ is an N-vector of holdings of state-contingent claims on exit from date 0, and $\mathbf{X}$ is the $N\times M$ asset payout matrix at date 1.

For all of the utility functions that we consider, with the exception of the reference-dependent preferences, the utility functional contains a weighted average of the utility of outcomes of the following form:

$$U(\mathbf{c};\mathbf{p})=\sum_i w_i(\mathbf{c},\mathbf{p})u(c_i).$$

In general, the dependence of weights on consumption distribution leads to nondifferentiable utility.[8] However, at the allocations where weights are locally constant, $U(\bullet)$ is differentiable and we can use the first-order necessary conditions to solve this problem. Differentiability is possible in all but a "small" set of allocations provided that the outcome function, $u(\cdot)$, is differentiable. For example, for DA preferences such allocations would be all but those on the DA threshold. For RDEU and AA preferences, these are all of the allocations with distinct consumption across all states.

After normalizing the price of consumption at time 0 to 1, we can express the rest of the state prices in a relatively standard way

$$\pi_i=w_i(\mathbf{c},\mathbf{p})\frac{u'(c_i)}{u'(c_0)}, i=1,\ldots,N. \tag{7}$$

The utility of outcomes function must be concave and differentiable, and the optimal consumption choices at date 1 must be distinct for all states, at least for some agents, in order to use these first-order conditions to solve the model. Our assumptions imply that sufficiently small perturbations to the consumption vector will not change the rankings of the outcomes. This means that the weighting vector can be treated as locally constant. Together with our first assumption, this implies that the value function will be locally concave and differentiable at the optimum.[9]

[8] For example, in the previously defined utility functions weights can depend on the rank of outcome or the location of outcome relative to the certainty equivalent-based threshold.

[9] These assumptions only make the first-order conditions necessary. Sufficiency would require, in addition, global concavity of the value function. Concavity of $u(\cdot)$ guarantees sufficiency for (G)DA and AA preferences because they are linear in probabilities. For RDEU, the above set of assumptions is insufficient to ensure global concavity and requires another restriction on the class of admissible weighting functions. These conditions and the proof of concavity of the RDEU are available on request.

We study the properties of asset prices in economies populated by heterogeneous agents. There are K agents, each with endowments e_0^k and $\mathbf{e}^{\mathrm{k}}$, solving the individual problem in (5) given asset prices π subject to individual budget constraints.In equilibrium, we also impose zero net supply of state contingent claims and the aggregate resource constraint

$$\sum_{k=1}^{K} \boldsymbol{\theta}^k = 0,\ \sum_{k=1}^{K} \mathbf{e_0^k} = \mathbf{e_0^a},\text{ and } \sum_{k=1}^{K} \mathbf{e^k} = \mathbf{e^a},$$

where $\mathbf{e}_0^{\mathrm{a}}$ and $\mathbf{e}^{\mathrm{a}}$ are the N-dimensional vectors of the aggregate endowments. We employ the standard notion of a competitive equilibrium where, given the prices of financial assets, the consumer types optimize their respective utilities and the markets for financial assets and the consumption good both clear.

The existence of the RA follows by standard arguments.[10] These model economies differ from a standard complete markets model with EU in two important respects. First, due to the first-order risk aversion it may be optimal for some agent types to insure consumption in some or in all states. This endogenous withdrawal from the market for risky assets means that a representative agent with state-independent weights is misspecified.[11] Second, in this equilibrium, it is not possible, in general, to characterize the solution using the gradient of the utility function because first-order risk averse preferences are not differentiable in all feasible allocations. In these cases, we rely on a numerical procedure to solve for equilibrium prices. This method is outlined in an Internet Appendix.[12]

III. Calibrating the Economies

In all of the model economies that we consider, there are $K = 2$ agent types. Again, this choice is made for simplicity, and all of our arguments carry over to economies with a larger number of agent types. The aggregate endowment at time 0 is normalized to 1. The states at date 1 are calibrated to match the unconditional distribution of gross annual real per capita consumption growth for the period from 1949 to 2006. This distribution is shown in Figure 2. We focus on this period because of the evidence in Chapman (2002) that shows that the mean and variance of consumption growth in the late nineteenth and the first half of the twentieth centuries are significantly larger than the mean

[10] The only nonstandard issue in proving the existence of the representative agent is to verify that we have made sufficient assumptions on preferences to ensure the global concavity of the value function. For (G)DA and AA utilities, concavity of utility of outcomes implies global concavity. For RDEU, in addition, a sufficient condition for global concavity requires a concave weighting transformation function. The proof of these conditions is available on request.

[11] The representative agent equilibrium for the case of RDEU, with full participation in risky asset markets, resembles equilibria with heterogeneous beliefs, as in Gollier (2003) and Jouini and Napp (2006). The primary difference between these types of models is that, with RDEU, the representative agent's beliefs are determined endogenously as a nonlinear function of the individual agents' decision weights.

[12] This Internet Appendix is available online in the "Supplements and Datasets" section at http://www.afajof.org/supplements.asp.

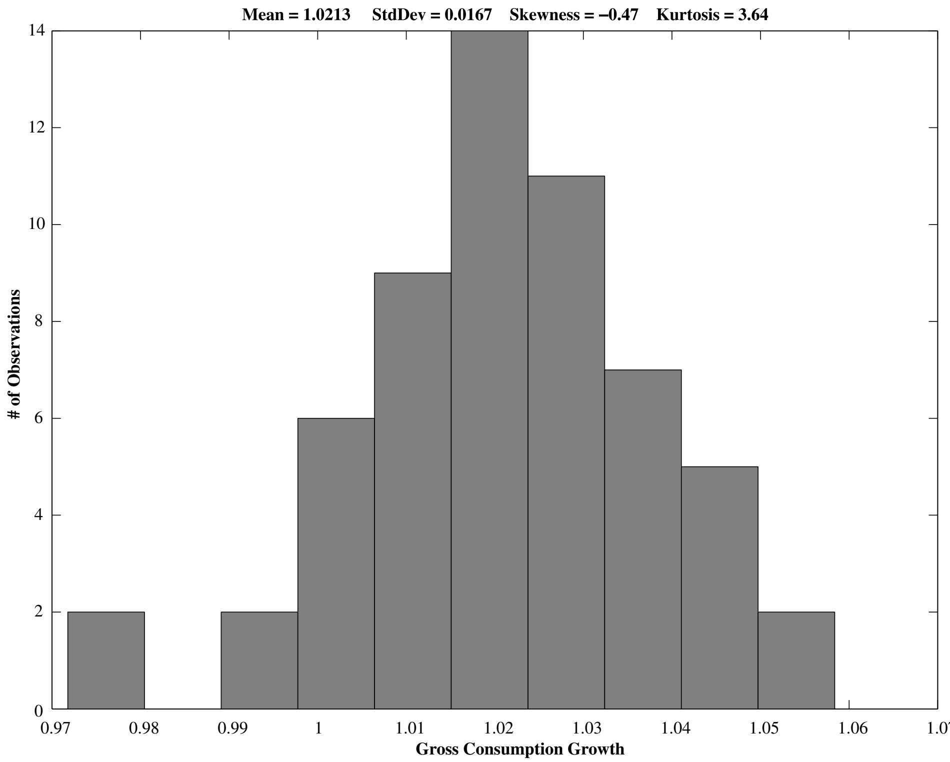

Figure 2. Gross consumption growth: 1949 to 2006. The histogram shows the empirical distribution of gross real per capita consumption growth measured. Consumption is total PCEs, and population is the total civilian noninstitutional population.

and variance of consumption growth over the latter half of the 20th century. We choose nine growth rate states at date 1 with each state corresponding to the mid-point of a bin on the histogram in Figure 2. The states and their associated probabilities are shown in Table I.

The first-order risk averse utility specifications examined in Section I can all be decomposed into a utility over outcomes function and a specification of either a decision weight function or a reference point. The utility over outcomes function that is most commonly used in the literature is power utility,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \tag{8}$$

for $\gamma > 0$ and $\gamma \neq 1$ and $u(c) = \log c$ for $\gamma = 1$, and we adopt this form of the function for each specification that we examine. In all of the examples, we set $\gamma = 2$.[13] This level is consistent with more conservative microeconomic-based calibrations of utility function curvature, although it should be noted

[13] To include time preference, the utility at time 1 is normally written in the form $\beta u(c)$, where β is the time discount parameter. For simplicity—and because it does not affect the main results—we set $\beta = 1$ throughout the paper.

Table I

Gross Endowment Growth Rate States

These endowment growth rates and their associated probabilities are calibrated to match the histogram in Figure 2.

State	Endowment	Probability
1	0.976	0.03
2	0.993	0.03
3	1.002	0.10
4	1.011	0.16
5	1.019	0.24
6	1.028	0.19
7	1.037	0.13
8	1.045	0.09
9	1.054	0.03

that standard estimates of utility function curvature are based on an expected utility specification. It is far too low relative to the parameter values that are required to rationalize the equity premium with a representative agent and aggregate consumption.

In order to understand the importance of preference-based heterogeneity for asset returns, for each utility specification we compare an economy with two equally wealthy consumers with different first-order risk aversion parameterizations to an economy with a single agent type whose utility function is a simple average of the preference parameters of the two consumer types. In an EU setting, the pricing implications of these two structures are, by construction, identical. Therefore, any differences that we observe between these two economies follow directly from the manner in which the preferences aggregate to the true representative agent.

There are no definitive studies that allow for specific calibration of the preferences defined in Section I. Therefore, for each specific choice, we choose a range of values that seems a priori reasonable and that includes the parameter values used by other researchers in prior examinations of these different utility functions. In the RDEU case, we consider values of the pessimism parameter, ϕ, that range from 0.4 to 1.0 (the EU case), which includes the values of $\{0.5, 0.75, 0.9\}$ used by Epstein and Zin (1990). In GDA, we use $\delta = 0.9692$ (the same value as Routledge and Zin (2004)), and we consider values of $\theta \in \{0, 9, 18, 27\}$, which includes the values of 9 and 24 used in Routledge and Zin (2004).

For ambiguity averse (AA) preferences, we parameterize ambiguous probability sets by a single parameter $\epsilon_i \geq 0$. The probability of each outcome, $\tilde{p}_j$, can vary around the objective probabilities, p_j, so that

$$\Pi_i = \left\{ \tilde{p} = [\tilde{p}_1, \ldots, \tilde{p}_n]' : \left| \frac{\tilde{p}_j}{p_j} - 1 \right| \leq \epsilon_i, \right.$$
$$\left. \sum_j \tilde{p}_j = 1, \text{ and } 0 \leq \tilde{p}_j \leq 1, \forall j = 1, \ldots, n \right\}.$$

When $\epsilon_i = 0$, Π_i is a singleton and coincides with the objective probability distribution. In this case the agent's preferences are EU. When $\epsilon_i > 0$, agents will choose the probability distribution from Π_i that, for a given consumption allocation, will minimize their expected utility. We experiment with $\epsilon_i \in [0, 0.95]$. A higher value of ϵ_i implies that the agent can have the subjective belief that any given state is unlikely, so that his subjective probability can be as low as 5% of the objective probability value.

IV. Results from the Calibrated Economies

We price the risk-free asset, r_f, and the market portfolio, r, which is the claim to aggregate consumption at date 1. Table II contains basic moments for the equity premium and the risk-free rate in the actual data and in a variety of first-order risk averse economies populated by homogeneous agents. The market return is measured as the annual value-weighted return to the New York Stock Exchange (NYSE)/Amex index from the Center for Research in Security Prices

Table II

Asset Return Moments in the Different First-Order Risk Aversion Models

The market return is measured as the annual value-weighted return to the NYSE/Amex index from the Center for Research in Security Prices (CRSP). Market returns are measured on the last trading day of the year. The risk-free rate is the yield on the 1-year discount bond from the Fama-Bliss data set at CRSP. Yields are measured on the last trading day of the year prior to the market returns. All nominal returns are converted to ex post real returns using the PCE deflator for nondurables and services. For all utility functions, the curvature of the utility of outcomes function is set to $\gamma = 2$, and the discount factor is set at $\beta = 1$.

Data	$E(r) - r_f$	$\sigma(r)$	r_f	Sharpe Ratio
1953–2006	7.273	17.149	1.181	0.424
Model	$E(r) - r_f$	$\sigma(r)$	r_f	Sharpe Ratio
RDEU				
$\phi_1 = \phi_2 = 1.0$	0.05	1.66	4.26	0.03
$\phi_1 = \phi_2 = 0.8$	0.42	1.65	3.54	0.25
$\phi_1 = \phi_2 = 0.6$	0.92	1.65	2.53	0.56
$\phi_1 = \phi_2 = 0.4$	1.65	1.63	1.06	1.01
$\phi_1 = \phi_2 = 0.2$	2.75	1.62	−1.19	1.70
GDA				
$\theta_1 = \theta_2 = 0$	0.05	1.66	4.26	0.03
$\theta_1 = \theta_2 = 9$	1.09	1.65	2.89	0.66
$\theta_1 = \theta_2 = 18$	1.74	1.65	2.02	1.06
$\theta_1 = \theta_2 = 27$	2.19	1.65	1.43	1.33
Ambiguity Averse				
$\varepsilon_1 = \varepsilon_2 = 0.00$	0.05	1.66	4.26	0.03
$\varepsilon_1 = \varepsilon_2 = 0.25$	0.37	1.65	3.63	0.22
$\varepsilon_1 = \varepsilon_2 = 0.50$	0.68	1.65	3.00	0.41
$\varepsilon_1 = \varepsilon_2 = 0.75$	0.99	1.64	2.38	0.60
$\varepsilon_1 = \varepsilon_2 = 0.95$	1.23	1.64	1.89	0.75

(CRSP) at the University of Chicago. The risk-free rate is measured as the yield-to-maturity on the 1-year discount bond from the Fama-Bliss data set on CRSP. All returns are converted from nominal to (ex post) real using the personal consumption expenditure (PCE) deflator for nondurables and services. The first line of Table II shows standard results on the magnitude of the equity premium, the level of an estimate of the real risk-free rate, real equity volatility, and the Sharpe ratio. They are reported here for ease of comparison with the model-generated data.

The parameterizations of the simple RDEU model with homogenous agents is, at best, a partial resolution of the equity premium puzzle. As the first line of the RDEU section of Table II shows, the level of consumption uncertainty is small when filtered through the lens of expected utility. The risk premium is only 11 basis points, and the risk-free rate is 4.26% per year. The magnitude of the equity premium increases with increased pessimism on the part of the agents in the economy, but the maximum value in the range of parameters that we consider is still only 2.75% per year.[14] This finding is consistent with Epstein and Zin (1990). The risk-free rate is monotonically declining with increased pessimism, and the real rate is actually negative for the largest pessimism parameter of $\phi = 0.2$.

There is no relation between pessimism and the volatility of the risky asset. It is (roughly) constant across all parameterizations of all utility functions at 1.6% per year. This follows from the simple two-period structure of our model economies. At date 1, the price of the market endowment is zero so the return to the risky asset is $r = D/P_0$, and $\sigma(r) = \sigma(D)/P_0$. So, asset volatility is (essentially) the volatility of the dividend yield, not the return. Since $\sigma(D)$ is calibrated exogenously the only impact on volatility is through the (small) changes in P_0 across parameterizations.

The pattern of asset returns for the homogeneous GDA (and AA) economies is qualitatively similar to that of the RDEU economy. The risk premium is increasing in DA (and ambiguity aversion), and the risk-free rate is declining in DA (and ambiguity aversion). The Sharpe ratio is also increasing in both of these economies because of the increase in the risk premium and the constant volatility across the parameterizations. The maximum risk premium in the GDA parameterizations is 2.19% per year (1.23% in the AA parameterizations), and the risk-free rate in the GDA economies ranges from 4.26% per year down to 1.43% per year (3.63% to 1.89% in the AA economies).

A. Rank-Dependent Expected Utility

Table III provides our first evidence on the impact of preference heterogeneity on asset returns with first-order risk averse preferences. It shows the differences in the risk premium and the risk-free rate between an economy with two

[14] It is harder to resolve the equity premium puzzle using only consumption data from the second half of the 20th century. Our estimate of the risk premium would be significantly larger if we included the more volatile data from the first half of the 20th century.

Table III

Asset Return Differences Across RDEU Economies

The table shows the differences in asset return moments between an economy with two types of RDEU investors and a single-consumer economy where the consumer has a preference parameter that is equal to the wealth-weighted average of the two consumers. Each entry in the table includes the difference (heterogeneous minus homogeneous) and the *absolute* difference as a proportion of the level in the two-consumer economy (shown in square brackets). The time preference parameter is set equal to 1.0, and the curvature parameter for the utility of outcomes function is set equal to 2.0. The special case of EU corresponds to $\phi = 1.0$.

	ϕ for Consumer 2			
ϕ for Consumer 1	0.4	0.6	0.8	1.0
Panel A: Expected Risk Premium (%)				
0.4	0			
0.6	−0.25	0		
	[25.02]			
0.8	−0.44	−0.17	0	
	[91.06]	[35.07]		
1.0	−0.54	−0.31	−0.12	0
	[497.4]	[287.4]	[111.0]	
Panel B: Risk-Free Rate (%)				
0.4	0			
0.6	0.50	0		
	[21.0]			
0.8	0.78	0.20	0	
	[23.5]	[6.1]		
1.0	1.10	0.47	0.21	0
	[26.3]	[11.7]	[5.1]	

types of RDEU investors and a single-consumer economy where the consumer has a preference parameter that is equal to the wealth-weighted average of the two consumers.[15] Each entry in the table includes the difference between the two-agent and the single-agent economies and the absolute difference as a proportion of the level in the two-agent economy (shown in square brackets). The time preference parameter is set equal to 1.0, and the curvature parameter for the utility of outcomes function, for both agents, is set equal to 2.0. The special case of EU corresponds to an agent with $\phi = 1.0$.

When the consumers are identical, there is no difference from the pricing in the average-consumer economy. So, the entries on the main diagonal in each panel are zero, by construction. When consumers are heterogeneous, identical agent economies overstate the equity premium compared to the heterogeneous agent economies. In heterogeneous agent economies the risk premium declines

[15] We do not report results for the volatility of the risky asset, since they are insensitive to variation in the preference parameters (see the discussion in the previous subsection).

and the difference depends on the wedge between parameters ϕ. The risk premium falls by as much as a factor of five when the difference in these parameters is the highest ($\phi = 1$ and $\phi = 0.4$), and the decline in risk premium is still significant even for moderate differences across agents shown in the entries closer to diagonal. The differences between the homogeneous and heterogeneous agent economies are larger as the two agents exhibit higher pessimism because these economies exhibit a larger scope for generating risk premiums (see Table II).

The differences in the risk-free rate between the two-agent and the average agent economy are shown in Panel B of Table III. The identical agent economies have lower risk-free rates than heterogeneous agent economies. Again, the differences can be substantial, although they are generally lower (in percentage terms) than the risk premium differences. This smaller impact of heterogeneity on the price of the risk-free asset is intuitively linked to the nature of risk sharing. In a heterogenous consumer economy, both consumers determine the price of the risk-free asset but the aggregate risk is priced only by one of the consumers. In the identical consumers economy, both consumers participate in pricing all assets. As a result, when one of the consumers disengages from the risk premium determination, we observe that heterogeneity has more impact on the risk premium than on the risk-free rate.

B. Disappointment Aversion

Table IV reproduces the results in Table III for the case of generalized disappointment aversion. As before, in the identical agent cases, the disappointment aversion is the average of the two agents' disappointment aversion coefficients. The differences across the different parameterizations of this economy are a function of the wedge between the disappointment aversion parameters. These differences are substantial when one of the agents has expected utility. In these cases, relying on a single-agent model overstates the risk premium by 10 to 25 times the level in the corresponding two-agent economy; see the first column of Panel A of Table IV. As in the RDEU economy, the homogeneous agent GDA economy overstates the risk premium (see the negative differences in Panel A of both Tables III and IV).

Panel B of Table IV shows that the risk-free rates are understated in the single-agent GDA economies, relative to the two-agent economies. This is qualitatively similar to the results in Table III, although the differences are increasing as we move away from an economy with a single EU agent type (compare the first column in Panel B to the other columns in Panel B). These differences are larger in both absolute value and as a percentage of the heterogeneous agent levels for the GDA economies compared to the RDEU economies. We also experiment with the DA model of Gul, which is a special case of GDA with $\delta = 1$. The results are qualitatively similar to the GDA cases reported in Table IV.[16]

[16] These results are available on request. In the case of the DA model, the more disappointment averse agent holds riskless consumption, similar to RDEU, because the kink in the indifference curves is located on the certainty line.

Table IV
Asset Return Differences Across GDA Economies

The table shows the differences in asset return moments between an economy with two types of GDA investors and a single-consumer economy where the consumer has a preference parameter that is equal to the wealth-weighted average of the two consumers. Each entry in the table includes the difference (heterogeneous minus homogeneous) and the *absolute* difference as a proportion of the level in the two-consumer economy (shown in square brackets). The time preference parameter is set equal to 1.0, the curvature parameter for the utility of outcomes function is set equal to 2.0, and the disappointment threshold is set at 0.97.

	θ for Consumer 2			
θ for Consumer 1	0	9	18	27
Panel A: Expected Risk Premium (%)				
0	0			
9	−0.58	0		
	[1010.7]			
18	−1.04	−0.34	0	
	[1794.5]	[30.5]		
27	−1.40	−0.63	−0, 22	0
	[2416.5]	[56.7]	[13.0]	
Panel B: Risk-Free Rate (%)				
0	0			
9	0.84	0		
	[19.4]			
18	1.38	−0.11	0	
	[32.1]	[4.8]		
27	1.83	0.26	−0.82	0
	[42.9]	[11.1]	[89.0]	

C. Ambiguity Aversion

The results for ambiguity aversion (AA) preferences are shown in Table V. Setting ε equal to zero corresponds to the case of no ambiguity. In this special case, the agent has EU preferences. In general, the table shows results that are consistent with our findings for both RDEU and GDA. The expected risk premium for different parameterizations of the two-agent model are shown in Panel A. As in RDEU, the single (weighted-average) agent model overstates the magnitude of the equity premium. The extent of this overstatement is always increasing in the difference in ambiguity between the two agents, with the largest difference (in both absolute and percentage terms) occurring for the case where Agent 2 is EU and Agent 1 has extreme ambiguity. These differences are not large in magnitude. However, given the small risk premiums in the homogeneous agent AA economies (Table II), these differences are large in percentage terms. For example, the first column in Panel A of Table V shows that the weighted-average single-agent economy has a risk premium that ranges from more than twice as high to almost nine times as large as the underlying two-agent model.

Table V

Asset Return Differences Across AA Economies

The table shows the differences in asset return moments between an economy with two types of AA investors and a single-consumer economy where the consumer has a preference parameter that is equal to the wealth-weighted average of the two consumers. Each entry in the table includes the difference (heterogeneous minus homogeneous) and the *absolute* difference as a proportion of the level in the two-consumer economy (shown in square brackets). The time preference parameter is set equal to 1.0, and the curvature parameter for the utility of outcomes function is set equal to 2.0.

	ε for Consumer 2				
ε for Consumer 1	0.00	0.25	0.50	0.75	0.95
Panel A: Expected Risk Premium (%)					
0.00	0				
0.25	−0.15	0			
	[238.2]				
0.50	−0.30	−0.15	0		
	[472.1]	[38.9]			
0.75	−0.46	−0.30	−0.12	0	
	[690.2]	[79.0]	[17.3]		
0.95	−0.58	−0.42	−0.27	−0.13	0
0.95	[854.1]	[110.6]	[39.5]	[10.7]	0
Panel B: Risk-Free Rate (%)					
0.00	0				
0.25	0.22	0			
	[5.32]				
0.50	0.53	−0.13	0		
	[12.76]	[3.96]			
0.75	0.84	[3.56]	−0.60	0	
	[12.76]	[3.56]	[27.63]		
0.95	1.09	0.31	−0.42	−0.97	0
	[26.15]	[10.06]	[20.69]	[82.24]	

The behavior of the risk-free rate in the AA economies is more complicated than in the RDEU and GDA models. When one of the agents is EU (see the first column in Panel B of Table V), the risk-free rate is overstated in the single-agent model, with the extent of the overstatement of the rate ranging from 5.3% to 26.2% of its value in the corresponding two-agent model.

This effect is weaker when both agents have ambiguity aversion. Initially, for higher levels of ambiguity, heterogeneity results in lower levels of the risk-free rate. As we move down the columns in Table V, the two-agent economy has ambiguity parameters that are equal to $\varepsilon_1 = \varepsilon + \Delta$ and $\varepsilon_2 = \varepsilon$, where ε is the value of the parameter in the diagonal cell and Δ is the increase in ambiguity for agent 1. In the corresponding single-agent economy, the ambiguity parameter is the average of the two values, $\varepsilon + \Delta/2$. The interest rate then declines in both the heterogeneous and homogeneous economies because the agents (or one agent) become more AA. This drives up the price of risk-free consumption.

This change occurs more slowly in the single-agent economy because the change is only half the size of the two-agent change. At some point in the heterogeneous economy, the most risk averse agent insures all or almost all states, at which point the homogeneous economy catches up. For example, consider the case of $\varepsilon_2 = 0.25$. The first entry is negative and then it becomes positive.

V. Reference-Dependent Preferences

A. Definition

Reference-dependent preferences evaluate outcomes by comparing them to a pre-specified (reference) level. Some of the earliest applications were based on habit formation utility. More recently, researchers have incorporated loss aversion elements of the prospect theory of Kahneman and Tversky (1979) and Tversky and Kahneman (1992) into these models. Köszegi and Rabin (2005, 2007) propose a general reference-dependent utility model that incorporates the idea of loss aversion and combines it with a standard utility of outcomes formulation. The Köszegi and Rabin utility function consists of two parts: a standard expected utility function and a utility of loss/gain function, denoted μ. The latter is defined on utility differences of actual outcomes from the utilities of reference points. The loss/gain function is asymmetric around zero, reflecting a well-known experimental fact that people receive considerably more disutility from losses than they receive utility from similarly sized gains.

The value function for reference-dependent utility is

$$V(c) = \alpha E[u(c)] + (1-\alpha)E[\mu(u(c) - u(r))], \tag{9}$$

where α determines the relative weights of the EU and the reference-dependent parts. The function $\mu(\cdot)$ is an asymmetric loss/gain function, defined later, and *r* is the reference level of consumption. One of the more common formulations of the gain/loss function, μ, was proposed by Kahneman and Tversky (1979)[17]:

$$\mu(x) = \begin{cases} x^{\theta}, & \text{if } x \geq 0 \\ -\lambda(-x)^{\theta}, & \text{if } x < 0. \end{cases}$$

The reference levels in the utility function in (9) remain to be specified. Köszegi and Rabin propose reference levels that emerge as endogenous expectations from the decision-maker's prior experiences in similar situations. This approach cannot be used in the context of our simplified single-period problem. Instead, the reference levels for a given consumption allocation must be specified using only the information available to the consumer at the time of the trading decisions. We consider several choices for references consumption levels, but they all have similar qualitative and quantitative results.[18] In the following

[17] In general, $\mu(x)$ must satisfy the following three sets of conditions: (i) *monotonicity*: $\mu(0) = 0$, continuous, strictly increasing; (ii) *saturation*: twice continuously differentiable for $x \neq 0$, $\mu''(x) \leq 0$ for $x < 0$, and $\mu''(x) \geq 0$ for $x > 0$; (iii) *loss aversion*: $\mu(y) + \mu(-y) < \mu(z) + \mu(-z)$ for $z > y > 0$ and $\lim_{x \downarrow 0} \frac{\mu(x)}{\mu(-x)} = \lambda > 1$.

[18] These results are available on request.

sections, the reference level is based on the certainty equivalent for a given allocation computed from the expected utility portion of the value function

$$r(c) = u^{-1}(E[u(c)]). \tag{10}$$

The complicated structure of the reference point makes any analysis involving reference-dependent utility less general than the analysis of the axiomatic utilities we consider earlier. It is also not feasible to apply general optimum existence Lagrange theorems toward the reference-dependent functions due to convexities and potential nonmonotonicity near certainty.[19] Nevertheless, the prominence of different versions of reference-dependent preferences in the behavioral asset pricing literature (see, for example, Barberis, Huang, and Santos (2001)) argue for explicitly considering these preferences as well. We are able to illustrate, later, that our main intuition applies in this case as well by considering a simplified two-state static setting that can be solved analytically. The examples we consider illustrate that the main intuition discussed for other types of first-order risk averse utilities also applies in the case of a reference-dependent utility function.

A two-state example of indifference curves for the utility function with this reference level is shown in Figure 3.[20] Note that on a "large scale" the indifference curves resemble standard expected utility, except apparent convergence at different angles toward the riskless allocations on the 45-degree line. This is similar to the kinks at certainty in the utility specifications examined in Figure 2. There is also a nonconvex "tip" at certainty that is due to loss aversion: The small-scale gambles near certainty generate a sense of loss (as in our choice of reference function) and are worse than a guaranteed payoff equal to the minimum across the two states.[21] The kink at the certainty line generates first-order risk averse behavior. Sufficiently loss averse agents may prefer to avoid taking risks due to losses in some states. These investors may choose not to participate in sharing aggregate risk of the economy and hold risk-free portfolios. We consider such quantitative examples further.

B. *Calibrating Reference-Dependent Utility*

We begin by assuming a two-state endowment process that matches the first two moments of the annual consumption data in Figure 2. In particular, let

[19] Relatedly, there are potentially undesirable normative properties of these utility functions that may imply inconsistent (nontransitive) choice (see Gul and Pesendorfer (2006)).

[20] For this figure, we use the following parameterization: $p_1 = p_2 = 0.5, \theta = 0.5, \alpha = 0.5, \lambda = 2$, and $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma = 2$.

[21] This feature is due to local nonmonotonicity stemming from loss aversion. It is also possible to turn this conclusion around by redefining the reference point to be the minimum of the two outcomes rather than being between the outcomes as in our example. Then, the best outcome is treated as a gain and the worst generates no disutility. In this case the agent would be better off reducing her payoff in one of the states just to experience a sense of a gain in another state. This reversal of predictions and nonmonotonicity of the utility demonstrate sensitivity of the reference-dependent utility to the choice of economic environment and particularly to the reference level.

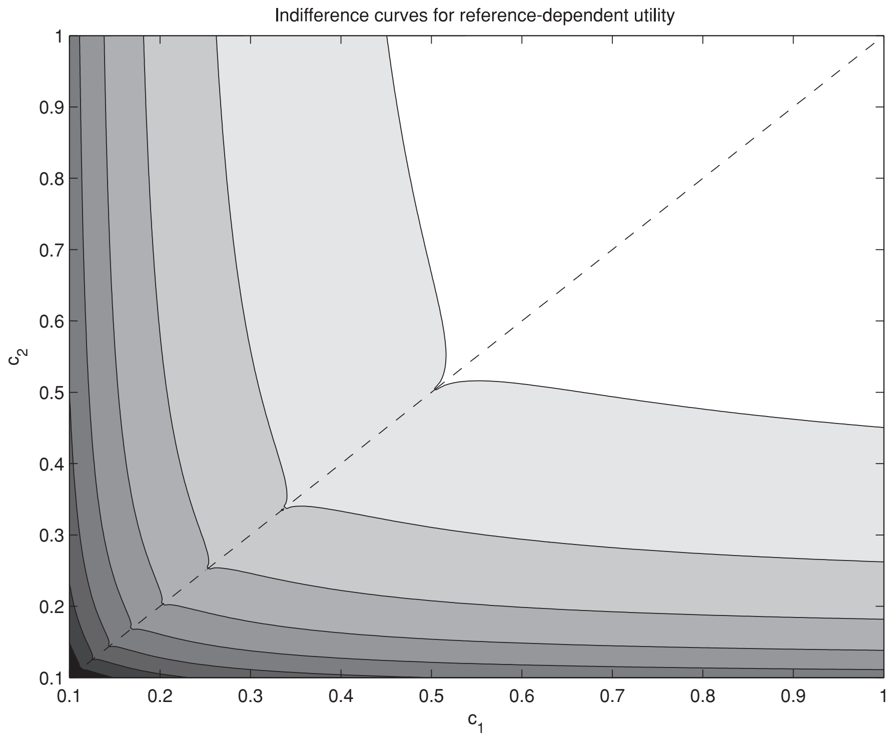

Figure 3. Indifference curves for a two-state version of reference-dependent utility. Consumption in the first state is c_1, and c_2 is consumption in the second state.

$e^a = [1.0046, 1.0380]$, where the two states are equally likely. There are three different calibrated economies: a homogeneous rank-dependent economy (with varying levels of α), a homogeneous EU economy, and an economy with one EU agent and one reference-dependent agent.

There is no guidance in the existing literature on the choice of the parameter that weights loss aversion versus standard utility, α. We focus on the value of $\alpha = 0.5$, a parameterization in which expected utility and reference-dependent agents are equally weighted. We also consider values of $\alpha \in \{0.25, 0.75\}$. The loss aversion parameter in our base parameterization is set to $\lambda = 2$. This value is used in some studies in the existing literature, but we also consider values of $\lambda \in \{1.5, 2.5\}$. Finally, the value of θ in our base parameterization is 0.8, but we also consider the values of 0.7 and 0.9.

C. Two-State, Two-Agent Equilibria with Reference-Dependent Utility

We consider parameterizations of equilibria where a reference-dependent agent interacts with an expected utility agent. We are particularly interested in the asset prices in equilibria where the reference-dependent agent holds a riskless consumption allocation.[22] The setup is modified from the one that we considered in the context of the other first-order risk averse utility functions as follows: There is no choice of consumption at time 0; that is, $u(c_0)$ and c_0 are omitted from the individual optimization problem defined in (5). In the absence of intertemporal choice, we normalize the price of the risk-free asset to 1. There are two states, 1 and 2, and consumption is indexed as c_i^k, where k and i index agents and states, respectively. This simpler problem can be treated analytically.

Let Agent 1 be the EU agent and Agent 2 be the reference-dependent agent. In equilibrium where only the EU agent bears risk, the state prices are determined by this agent's marginal utility, while the reference-dependent agent holds a riskless consumption allocation at the "kink" of his indifference curves. The equations that define equilibrium state price 1 (state price 2 is normalized to 1) are given by the first-order condition for Agent 1, the aggregate resource constraint, and the individual budget constraints as follows:

$$\begin{aligned} c_1^1 + c_1^2 &= e_1^a \\ c_1^2 + c_2^2 &= e_2^a \\ \pi_1 c_1^1 + c_2^1 &= \pi_1 e_1^1 + e_2^1 \\ \pi_1 c_2^1 + c_2^2 &= \pi_1 e_1^2 + e_2^2 \\ \pi_1 u'\big(c_2^1\big) &= u'\big(c_1^1\big). \end{aligned}$$

For the power utility function, $u(c) = c^{1-\gamma}/(1-\gamma)$, the last equation can be written as

$$\pi_1^{\frac{1}{\gamma}} c_2^1 = c_1^1.$$

These equations imply the following nonlinear equation for the state price

$$\big(\pi_1 e_1^1 + e_1^2\big)(1+\pi_1) + \big(\pi_1 e_1^2 + e_2^2\big)\Big(\pi_1^{\frac{1}{\gamma}} + \pi_1\Big) = e_1^a(1+\pi_1)\Big(\pi_1 + \pi_1^{\frac{1}{\gamma}}\Big). \tag{11}$$

In the parameterized examples later we use this equation to compute state price.

When we consider equilibria with two identical reference-dependent agents, the state price (when it is defined) is given by the marginal utility at the consumed allocation.[23] The derivative of the expected utility portion of the

[22] Due to the aforementioned complications with reference-dependent utility, we do not consider a mixture of two reference-dependent agents with different utility functions. The reason is that the equilibrium may fail to exist for some parameter configurations due to a discontinuity of demands in prices stemming from nonmonotonicity of the utility. A detailed analysis of this issue is outside the scope of the present paper.

[23] A potential problem is that there exist allocations near certainty where utility is nonmonotone and convex and where the demand functions may be discontinuous in prices. For such allocations there may not exist any price that would imply that the agent consumes his endowment. We checked that consumption endowments (optimal choice in case of identical agents) are on the concave part of the utility in our parameterizations so that the solution to the problem is well-defined.

reference-dependent functional (using the form of the budget constraint) is

$$\frac{\partial Eu(c)}{\partial c_1} = p_1 u'(c) - p_2 u'(c_2)\pi_1.$$

Differentiating the full utility functional, we find that

$$\begin{aligned}\frac{\partial V(c)}{\partial c_1} &= \alpha \frac{\partial Eu(c)}{\partial c_1} \\ &\quad + (1-\alpha)\left\{ p_1 \mu(c_1 - Eu(c)) \left(u'(c_1) - \frac{\partial Eu(c)}{\partial c_1} \right) \right. \\ &\quad \left. - p_2 \mu(c_1 - Eu(c)) \left(\pi_1 u'(c_2) + \frac{\partial Eu(c)}{\partial c_1} \right) \right\}.\end{aligned}$$

We obtain the expression for the state price, in the case of identical agents, by setting this derivative to zero and rearranging:

$$\pi_1 = \frac{[\alpha - (1-\alpha)E\mu'(c - Eu(c))]p_1 u'(c_1) + p_1 \mu'(c_1 - Eu(c))u'(c_1)(1-\alpha)}{[\alpha - (1-\alpha)E\mu'(c - Eu(c))]p_2 u'(c_2) + p_2 \mu'(c_2 - Eu(c))u'(c_2)(1-\alpha)}.$$

Note the similarity of this equation to the standard one for EU where the price is simply given by $\pi_1 = \frac{p_1 u'(c_1)}{p_2 u'(c_2)}$. The above formula collapses to the one for EU when $\alpha = 1$.

We now use the two-state, two-agent model to demonstrate the effects of preference heterogeneity with reference-dependent utility. We consider an equilibrium with two agents, one of which has reference-dependent utility and the other of which has standard expected utility. We focus on an economy where the loss averse agent optimally does not participate in aggregate risk sharing. The outcome of this economy is compared to two different homogenous-agent economies: one where the representative agent has expected utility and one where the representative agent has reference-dependent utility.

Table VI shows the value of the risk premium in the different reference-dependent utility economies. The benchmark parameterization is shown in the first row. In the homogeneous (single-agent) economy, the risk premium varies from a low of 0.49% per period to a high of 3.81% per period, and it has a value of 1.34% per period in our benchmark parameterization. The risk premium is higher when the reference-dependent agent places a higher weight on the non-EU component of preferences (see equation (9)), that is, when α is higher and when the nonlinearity of the reference-dependent component is greater, that is, when λ is higher. In contrast, the risk premium in the two-agent economy is constant at 0.11% per period. This follows because the EU agent provides complete consumption insurance to the reference-dependent agent and sets the level of the risk premium in *all* of the parameterizations considered in Table VI. The last two columns of Table VI show that, consistent with our findings in the multi-state economies for the other first-order risk averse preferences, the weighted-average-agent economy significantly overstates the equity premium that would hold in all of the parameterizations of the two-agent economy.

Table VI

Asset Returns in the Reference-Dependent Models

The table shows the value of the risk premium in different economies with equally weighted reference dependent and expected utility agents. α determines the weight that the reference-dependent agent places on the EU component versus the reference-dependent component of utility; see equation (9) in the text. λ determines the asymmetry in the loss/gain component of the reference dependent portion of utility, and θ determines the curvature of the loss/gain function. The risk-free rate is normalized to zero and the risk premium is reported for different parameter choices for the reference-dependent agent. The risk premium is reported in percent per period. The last two columns compute the difference—in level and in percent of the heterogeneous agent value—between the economy with a single reference dependent investor ("identical") and the two-agent economy ("heterogeneous").

				Risk Premium		Difference	
Case	α	λ	θ	Homogeneous	Heterogeneous	Level	Percent
Benchmark	**0.50**	**2.0**	**0.8**	**1.34**	**0.11**	**−1.23**	**1,153**
2	0.50	1.5	0.8	0.70	0.11	−0.59	555
3	0.50	2.5	0.8	1.97	0.11	−1.86	1,743
4	0.25	2.0	0.8	3.81	0.11	−3.71	3,469
5	0.75	2.0	0.8	0.49	0.11	−0.38	355
6	0.50	2.0	0.7	1.63	0.11	−1.53	1,431
7	0.50	2.0	0.9	1.08	0.11	−0.97	912

VI. Conclusions

The results of our evaluation of a wide variety of nonexpected utility preferences is that the expected risk premium and the risk-free rate are sensitive to ignoring even simple (two-agent type) heterogeneity. These are simple static models, and they still generate substantial differences in both risk premiums and risk-free rates between two-agent and single- (weighted average) agent versions of the same economy.

While the examples we consider are simple, they capture a basic intuition: *in general*, first-order risk aversion implies an unwillingness to assume any risk whenever the risk premium is deemed insufficiently high; see Segal and Spivak (1990) for a general proof of this result. Therefore, our findings extend well beyond the simple finite-state static environment of our examples: More disappointment averse/pessimistic/loss averse/ambiguity averse agents may choose not to bear *any* risk, which leaves the pricing of risk to the remaining agents. In these cases, it is incorrect to use single average agent pricing in lieu of the information contained in the entire cross-section of heterogeneous agents (and their wealth distribution). The risk premium in these economies will be influenced primarily by the agents who are most willing to bear risk, and it will differ considerably from the level implied by the "typical" preference parameters.

The participation mechanism that we highlight has the potential to generate time-varying effects on asset return moments as the wealth distribution changes over time and as different agents (endogenously) enter and leave the market. Indeed, there is a growing literature that examines the dynamics of single-agent models with first-order risk averse preferences.[24] The analysis of such dynamics requires a model that is substantially more complicated than the ones that we use, and introducing it here would obscure our fundamental point. We leave a thorough analysis of the dynamics of this class of models as an open area for future research.

REFERENCES

Anderson, Evan W., 2005, The dynamics of risk-sensitive allocations, *Journal of Economic Theory* 125, 93–150.

Ang, Andrew, Geert Bekaert, and Jun Liu, 2005, Why stocks may disappoint, *Journal of Financial Economics* 76, 471–508.

Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, The implications of first-order risk aversion for asset market risk premiums, *Journal of Monetary Economics* 40, 3–39.

Brandt, Michael W., Amir Yaron, and Qu Zeng, 2005, First order risk aversion, aggregation, and dynamic asset pricing, Working paper, Duke University, University of Pennsylvania, and University of Melbourne.

Chapman, David A., 2002, Does intrinsic habit formation actually resolve the equity premium puzzle? *Review of Economic Dynamics* 5, 618–645.

Constantinides, George M., 1982, Intertemporal asset pricing with heterogeneous consumers and without demand aggregation, *Journal of Business* 55, 253–267.

Dow, James, and Sergio Ribeiro da Costa Werlang, 1992, Uncertainty aversion, risk aversion, and the optimal choice of portfolio, *Econometrica* 60, 197–204.

Ellsberg, Daniel, 1961, Risk, ambiguity, and the savage axioms, *Quarterly Journal of Economics* 75, 643–669.

Epstein, Larry G., and Martin Schneider, 2006, Ambiguity, information quality, and asset pricing, *Journal of Finance* 63, 197–228.

Epstein, Larry G., and Stanley E. Zin, 1990, "First-order" risk aversion and the equity premium puzzle, *Journal of Monetary Economics* 26, 387–407.

Epstein, Larry G., and Stanley E. Zin, 2001, The independence axiom and asset returns, *Journal of Empirical Finance* 8, 537–572.

Gilboa, Itzhak, and David Schmeidler, 1989, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* 18, 141–153.

Gollier, Christian, 2003, Who should we believe? Collective risk-taking decisions with heterogeneous beliefs, IDEI Working Paper No. 201.

Gul, Faruk, 1991, A theory of disappointment aversion, *Econometrica* 59, 667–686.

Gul, Faruk, and Wolfgang Pesendorfer, 2006, The revealed preference implications of reference dependent preferences, Working paper, Princeton University.

Jouini, Elyes, and Clotilde Napp, 2006, Heterogeneous beliefs and asset pricing in discrete time: An analysis of pessimism and doubt, *Journal of Economic Dynamics and Control* 30, 1233–1260.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–292.

Köszegi, Botond, and Matthew Rabin, 2005, A model of reference-dependent preferences, *Quarterly Journal of Economics* 121, 1133–1165.

[24] Barberis, Huang, and Santos (2001) look at prospect theory, Routledge and Zin (2004) use generalized disappointment aversion, and Epstein and Schneider (2006) consider ambiguity aversion. Anderson (2005) examines some of the considerable technical issues with these dynamic models. Brandt, Yaron, and Zeng (2005) are the first (to our knowledge) to consider the asset pricing implications of heterogeneity with a form of dynamic first-order risk averse preferences.

Köszegi, Botond, and Matthew Rabin, 2007, Reference-dependent risk attitudes, *American Economic Review* 97, 1047–1073.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Polkovnichenko, Valery, 2005, Household portfolio diversification: A case for rank-dependent preferences, *Review of Financial Studies* 18, 1467–1502.

Quiggin, John, 1982, A theory of anticipated utility, *Journal of Economic Behavior and Organization* 3, 323–343.

Quiggin, John, 1993, *Generalized expected utility theory* (Kluwer Academic Publishers, Boston, MA).

Routledge, Bryan R., and Stanley E. Zin, 2004, Generalized disappointment aversion and asset prices, NBER Working Paper No. 10107.

Segal, Uzi, and Avia Spivak, 1990, First Order Versus Second Order Risk Aversion, *Journal of Economic Theory* 51, 111–125.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

Yaari, Menahem E., 1987, The dual theory of choice under risk, *Econometrica* 55, 95–115.