

Ambiguity in Asset Markets: Theory and Experiment*

Peter Bossaerts

California Institute of Technology, Swiss Finance Institute & CEPR

Paolo Ghirardato

Università di Torino & Collegio Carlo Alberto

Serena Guarneschelli

McKinsey, Inc.

William R. Zame

University of California, Los Angeles

First version: 03/24/03; This version: 07/07/07

*An earlier version circulated (March 2003-August 2004) under the title “The Impact of Ambiguity on Prices and Allocations in Competitive Financial Markets.” We are grateful for comments from audiences at Bocconi, CIRANO, Collegio Carlo Alberto, Columbia, East Anglia, Iowa, Irvine, Johns Hopkins, Kobe, Mannheim, Mainz, Northwestern, Nottingham, Sankt Gallen, the Securities and Exchange Commission, the 2003 Econometric Society North-American Summer Meeting (Evanston, IL), the XI Foundations of Utility and Risk Conference (Paris, July 2004), the 2005 Atlanta Fed Conference on Experimental Finance, the 2006 Skinance Conference (Hemsedal, Norway), and the 2006 European Summer Symposium on Financial Markets. Financial support was provided by the the Caltech Social and Information Sciences Laboratory (Bossaerts, Zame), the John Simon Guggenheim Foundation (Zame), the R. J. Jenkins Family Fund (Bossaerts), the Italian MIUR (Ghirardato), the National Science Foundation (Bossaerts, Zame), the Swiss Finance Institute (Bossaerts) and the UCLA Academic Senate Committee on Research (Zame). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or any other funding agency.

Abstract

This paper studies the impact of ambiguity and ambiguity aversion on equilibrium asset prices and portfolio holdings in competitive financial markets. It argues that attitudes toward ambiguity are heterogeneous across the population, just as attitudes toward risk are heterogeneous across the population, but that heterogeneity of attitudes toward ambiguity has different implications than heterogeneity of attitudes toward risk. In particular, when some state probabilities are not known, agents who are sufficiently ambiguity averse find *open* sets of prices for which they refuse to hold an ambiguous portfolio. This suggests a different cross-section of portfolio choices, a wider range of state price/probability ratios and different rankings of state price/probability ratios than would be predicted if state probabilities were known. Experiments confirm all of these suggestions. Our findings contradict the claim that investors who have cognitive biases do not affect prices because they are infra-marginal: ambiguity averse investors have an *indirect* effect on prices because they change the per-capita amount of risk that is to be shared among the marginal investors. Our experimental data also suggest a positive correlation between risk aversion and ambiguity aversion that might explain the “value effect” in historical data.

1 Introduction

The most familiar model of choice under uncertainty follows Savage (1954) in positing that agents maximize expected utility according to subjective priors. However, Knight (1939), Ellsberg (1961) and others argue that agents distinguish between risk (known probabilities) and ambiguity (unknown probabilities), and may display aversion to ambiguity, just as they display aversion to risk.¹ The financial literature, while admitting the possibility that some individuals might be averse to ambiguity, has largely ignored the implications for financial markets.²

In this paper, we use theory and experiment to study the effect of attitudes toward ambiguity on portfolio choices and asset prices in competitive financial markets. Our point of departure is the (theoretical) observation that aversion to ambiguity has different implications for choices — and hence, different implications for prices — than aversion to risk. Agents who are merely averse to risk will choose to hold a riskless portfolio (that is, a portfolio that yields identical wealth across all states) only if price ratios are exactly equal to ratios of expected payoffs, which is a knife-edge condition. However, agents who are averse to ambiguity will choose to hold an unambiguous portfolio (that is, a portfolio that yields identical wealth across states whose probabilities are not known) for an *open* set of prices and probabilities. If aversion to ambiguity is heterogeneous across the population and aggregate wealth differs across ambiguous states (states whose probability is not known), this generates a bi-modal distribution, with the most ambiguity averse agents holding equal wealth in ambiguous states and the other agents holding the net aggregate wealth. As a result, state price/probability ratios (ratios of state prices to probabilities) may be quite different than they would be if all agents maximized expected utility with respect to a common prior, even to the extent that pricing may be *inconsistent* with the preferences of a repre-

¹Knight used the terms risk and uncertainty; we use risk and ambiguity because they seem less likely to lead to confusion.

²Exceptions include Epstein & Wang (1994) and Cagetti, Hansen, Sargent & Williams (2002).

sentative agent who maximizes state-independent utility.

Our experimental findings confirm the predictions of this theoretical analysis. We find that a significant fraction of agents are highly ambiguity averse and refuse to hold an ambiguous portfolio, that ambiguity aversion is heterogeneous across the population, and that rankings of state price/probability ratios are anomalous exactly in those configurations when theory predicts they are most likely to be.

The environment we study is inspired by Ellsberg (1961). Uncertainty in Ellsberg’s environment is identified with the draw of a single ball from an urn that contains a known number of balls, of which one third are known to be red and the remainder are blue or green, in unknown proportions. Ellsberg asked subjects first, whether they would prefer to bet on the draw of a red ball or of a blue ball, or on the draw of a red ball or a green ball, and second, whether they would prefer to bet on the draw of a red or green ball, or of a blue or green ball. Ellsberg found (and later experimenters have confirmed) that many subjects prefer “red” in each of the the former choices and “blue or green” in the latter. Such behavior is “paradoxical” — that is, inconsistent with maximizing expected utility with respect to any subjective prior, and hence violates the Savage (1954) axioms, specifically, the “sure thing principle.”

We imbed this environment in an asset market in which Arrow securities (assets) are traded. Each security pays a fixed amount according to the color of the ball drawn from an Ellsberg urn. The Red security (i.e., the security that pays when a red ball is drawn) is risky (the distribution of its payoffs is known) while the Blue and Green securities are ambiguous (the distribution of their payoffs is unknown). In order to study the effects of ambiguity aversion, we exploit the freedom of the laboratory setting to augment the environment in three ways: first, by determining aggregate supplies of the various securities we manipulate aggregate wealth in the various states; second, by determining the number of balls of each color and by drawing balls without replacement, we manipulate true probabilities; third, by replicating sessions, we construct environments which are parallel in every dimension

except that in one environment the true composition of the urn is known and in the other environment it is unknown.^{3,4}

To model preferences that display ambiguity aversion, we use the multiple prior “ α -maxmin” model of Ghirardato, Maccheroni & Marinacci (2004), which is a generalization of the “maxmin” model of Gilboa & Schmeidler (1989). This specification provides a natural way to broaden the spectrum of agents’ behavioral traits, without a radical departure from the familiar expected utility model and with little loss in terms of tractability. For these preferences and experimental environment, the parameter α corresponds to the degree of ambiguity aversion: $\alpha = 1$ corresponds to extreme ambiguity aversion, $\alpha = 1/2$ corresponds to ambiguity neutrality, and $\alpha = 0$ corresponds to extreme ambiguity loving.

Ambiguity aversion ($\alpha > 1/2$) has implications for individual choice behavior: there is an *open* set of prices with the property that an ambiguity averse agent who faces these prices will always choose to hold an unambiguous portfolio (in our setting, a portfolio yielding equal wealth in the Green and Blue states). Indeed, an agent who is maximally ambiguity averse ($\alpha = 1$) will *always* choose to hold an unambiguous portfolio, no matter the relative prices of the ambiguous securities. By contrast, an agent who maximizes expected utility with respect to a subjective prior will choose to hold equal quantities of two securities only if the ratio of prices is *equal* to the ratio of subjective probabilities.

In a market in which attitudes toward ambiguity are heterogeneous across

³The behavior seen in Ellsberg’s paradox might suggest that the price of the Red security should be higher than the price of the Blue security and of the Green security, and that the price of the portfolio consisting of one Blue and one Green security should be higher than the price of the portfolio consisting of one Red and one Blue security. However, such prices could not obtain at a market equilibrium because they admit an arbitrage opportunity.

⁴Epstein & Miao (2003) studies an environment in which agents are equally ambiguity averse but have different information, and hence do not agree on which states are ambiguous. In our environment, agents agree on which states are ambiguous but exhibit differing levels of ambiguity aversion.

the population and the supplies of ambiguous securities (Blue and Green, in our case) are different, this choice behavior has an obvious implication for the cross-section of equilibrium portfolio holdings. Because the most ambiguity averse agents hold an unambiguous portfolio, less ambiguity averse agents must hold the imbalance of ambiguous securities. Thus, the cross-section of portfolio holdings should have a different mode and higher variation when there are ambiguity averse agents than when all agents maximize expected utility.

More subtly, ambiguity aversion also has implications for equilibrium pricing. If all agents maximize expected utility with respect to common priors (but with possibly different risk attitudes), equilibrium state price/probability ratios will be ranked oppositely to aggregate wealth and equilibrium prices can always be rationalized by a representative agent who maximizes expected utility (with respect to the common prior).⁵ Things change if some agents are ambiguity averse: Agents who are very ambiguity averse will refuse to hold an ambiguous portfolio. If the ambiguous securities are in unequal total supply, this means that the remaining agents — who are less ambiguity averse or who maximize expected utility with respect to a subjective prior — must hold the imbalance of ambiguous securities. This may distort state price/probability ratios; if the distortion is sufficiently large, state price/probability ratios may not be ranked opposite to aggregate wealth, in which case equilibrium prices *cannot* be rationalized by a representative agent who maximizes expected utility or even by a representative agent who is ambiguity-averse, at least under common “beliefs.”⁶ This would seem to have important implications for finance, where the representative agent methodology is pervasive.

Our laboratory environment is ideal for studying these predictions. We obtain a complete record of individual portfolio choices. We can manipulate supplies of ambiguous securities so that anomalous orderings are (predicted to be) likely in some treatments and unlikely in others. And we can compare outcomes in a treatment where some states are ambiguous with outcomes in

⁵See Constantinides (1982), for example.

⁶We use state price/probability ratios computed from a uniform prior over ambiguous states; we discuss alternative notions later in the paper.

a treatment which is identical in every respect *except* that state probabilities are commonly known.

Our experimental data are consistent with the theoretical predictions. The population is heterogeneous: some agents are quite ambiguity averse and some are not. In treatments where there is no ambiguity, the cross section of portfolio weights shows a single mode equal to the market weight; that is, the modal investor holds the market portfolio.⁷ In treatments where there is ambiguity, the mode is at equal weighting, reflecting the desire of highly ambiguity averse agents to hold ambiguous state securities in exactly equal proportions. (There may be a second mode at the net market weighting.) In treatments where there is no ambiguity, the ranking of state price/probabilities is opposite the the ranking of aggregate wealth; in treatments where there is ambiguity, the rankings are anomalous exactly in those treatments where theory predicts anomalous rankings are most likely.

One other feature of our experimental data is worth noting. In principle, there need be no correlation between ambiguity aversion (measured by α) and risk aversion (measured by concavity of u), but our experimental data suggests that a positive correlation may in fact obtain. If this is a property of the population as a whole, it could have significant effects on the pricing of different kinds of assets, and presents a potential explanation of the “value effect” — the observation that the historical average return of growth stocks is smaller than that of value stocks, even after accounting for risk. To the extent that growth stocks can be associated with ambiguity and value stocks can be associated with risk, heterogeneity in ambiguity aversion and positive correlation between ambiguity aversion and risk aversion would suggest that the markets for growth and value stocks should be segmented, and that growth stocks should be held — and priced — primarily by investors who are less ambiguity averse and hence (because of the presumed correlation) less risk averse, while value stocks would be held and priced by the market as a whole. This would suggest that growth stocks should carry a lower risk

⁷Some models — CAPM for instance — would predict that *all* agents should hold the market portfolio; the data do not support that prediction.

premium and yield lower returns, while value stocks should carry a higher risk premium and yield higher returns. As noted, this precisely what the historical data suggest; see Fama & French (1998) for instance.⁸

The approach here follows Bossaerts, Plott & Zame (2007), who study environments with pure risk (i.e., known probabilities). Bossaerts, Plott & Zame (2007) document that there is substantial heterogeneity in preferences but that much of this heterogeneity washes out in the aggregate, so that the pricing predicted by familiar theories such as CAPM (approximately) obtains even though portfolio separation does not. In the environment addressed here, with both risk and ambiguity, heterogeneity does not wash out in the aggregate and pricing predicted by familiar theories does not obtain.

As do we here, Easley & O’Hara (2005) also point out that the risk premium in markets populated with investors with heterogeneous attitudes towards ambiguity will depend on the number of investors who choose to hold aggregate risk, and derives (theoretical) implications for regulation, under the assumption that risk aversion and ambiguity aversion are uncorrelated. Unlike the present paper, Easley & O’Hara provide no empirical analysis to suggest that their assumptions about risk aversion and ambiguity aversion or their theoretical predictions are actually observed.

The present paper adds to an emerging literature that uses experimental evidence to motivate studying the effects of “irrational” preferences (eg., preferences other than expected utility) on prices and choices in competitive markets through experiments. Gneezy, Kapteyn & Potters (2003) analyze the impact of myopic loss aversion on pricing, but assumes homogeneous preferences. Kluger & Wyatt (2004) study the impact of particular cognitive biases on updating and pricing in experimental markets, but does not provide a theoretical framework within which it is possible to understand the effects (if any) of heterogeneity. Chapman & Polkovnichenko (2005) study the effects of a particular class (rank-dependent-utility) of “irrational” preferences on asset prices and portfolio holdings, but the preferences studied do not display ambiguity aversion in the sense studied here and equilibrium prices always

⁸We thank Nick Barberis for this observation.

admit a representative agent rationalization. See Fehr & Tyran (2005) for an overview.

A related literature, including Epstein & Wang (1994), Uppal & Wang (2003), Cagetti, Hansen, Sargent & Williams (2002), Maenhout (2000), Skiadas (2005), Trojani, Leippold & Vanini (2005), seeks to explain the equity premium puzzle (high average returns on equity and low average riskfree rate) by appealing to ambiguity (which they call Knightean or model uncertainty) on the basis of a model with an ambiguity-averse representative agent. Because ambiguity aversion does not seem to aggregate across a heterogeneous population, so that prices may not be rationalizable by any representative agent, our finding that there is substantial heterogeneity would seem to suggest problems with this literature.

Following this Introduction, Section 2 begins by presenting the theoretical analysis, generating predictions about choices and prices. Section 3 describes our experimental design. Section 4 analyzes the data in view of the theoretical predictions. Section 5 explores alternative explanations for the observed patterns in prices and holdings. Section 6 concludes.

2 Theory

We treat a market that unfolds over two dates, with uncertainty about the state of nature at the second date. In keeping with the Ellsberg experiment, we refer to the three possible states of nature as Red, Green, Blue or R, G, B .⁹ Trade takes place only at date 0; consumption takes place only at date 1. There is a single consumption good.

At date 0, each of N agents are endowed with and trade a riskless asset (cash) and Arrow securities whose payoffs depend on the realized state of nature. It is convenient to denote the security by the state in which it pays; thus the Red security pays 1 unit of consumption if the realized state is Red and nothing in the other states, etc. Write $p = (p_R, p_G, p_B)$ for the vector of prices of Arrow securities. Normalize so that the price of the riskless security is 1; absence of arbitrage implies that $p_R + p_G + p_B = 1$. Because a complete set of Arrow securities are traded, markets for contingent claims are complete (the riskless asset is redundant), so it is convenient to view our market as an Arrow-Debreu market for complete contingent claims.

Agents are completely described by consumption sets, which we take to be \mathbb{R}^3 , endowments $e \in \mathbb{R}^3$, and utility functions $U: \mathbb{R}^3 \rightarrow \mathbb{R}$. (To be consistent with the experimental environment described in Section 3 we allow wealth to be negative in some states.) An agent whose endowment is e and utility function is U and who faces prices $p \in \mathbb{R}_+^3$, chooses wealth $w \in \mathbb{R}^3$ to maximize $U(w)$ subject to the budget constraint $p \cdot w \leq p \cdot e$.

As usual, an *equilibrium* consists of prices p and individual choices w^n such that

- agent n 's choice w^n maximizes utility $U^n(w^n)$ subject to the budget constraint $p \cdot w^n \leq p \cdot e^n$

⁹Obviously the choice of labels is arbitrary; we maintain the Ellsberg labeling for ease of reference. In the experiments, we use the more neutral labeling X, Y, Z .

- the market clears:

$$\sum_{n=1}^N w^n = \sum_{n=1}^N e^n = W$$

2.1 Individual Choice: Expected Utility

We first recall familiar implications of the assumption of expected utility for choice behavior.

Consider an agent who maximizes expected utility according to (objective or subjective) priors π_R, π_G, π_B . By definition, this means the agent's utility for state-dependent wealth w is

$$U(w) = \pi_R u(w_R) + \pi_G u(w_G) + \pi_B u(w_B)$$

where u is felicity for certain consumption, assumed to be twice differentiable, strictly increasing and strictly concave. Given prices $p = (p_R, p_G, p_B)$ (and recalling that we allow wealth to be negative) the first order conditions for optimality are that

$$\frac{\pi_\sigma u'(w_\sigma)}{p_\sigma} = \frac{\pi_\nu u'(w_\nu)}{p_\nu} \text{ for all states } \sigma, \nu = R, G, B \quad (1)$$

Strict concavity implies that u' is a strictly decreasing function, so that $u'(w_\sigma) < u'(w_\nu)$ exactly when $w_\sigma > w_\nu$. Hence choices of state-dependent wealth are ranked oppositely to state price/probability ratios:

$$w_\sigma > w_\nu \iff \frac{p_\sigma}{\pi_\sigma} < \frac{p_\nu}{\pi_\nu} \text{ for all states } \sigma, \nu = R, G, B \quad (2)$$

(Note that the *ranking* of state-dependent wealth choices is independent of the felicity function u and of the magnitudes of prices, but of course the *magnitude* of wealth choices depends on both u and the the magnitude of prices.)

2.2 Individual Choice: Ambiguity Aversion

As we show, the implications of the assumption of ambiguity aversion for choice behavior may be quite different from those derived above.

As in the Ellsberg environment, we assume the true probability π_R is known but that π_G, π_B are unknown. To model ambiguity aversion, we follow Ghirardato, Maccheroni & Marinacci (2004) in assuming that utility is of the α -maxmin form

$$\begin{aligned}
 U(w) = & \pi_R u(w_R) \\
 & + \alpha \min_{\beta \in [0, (1-\pi_R)]} [\beta u(w_G) + (1 - \pi_R - \beta)u(w_B)] \\
 & + (1 - \alpha) \max_{\gamma \in [0, (1-\pi_R)]} [\gamma u(w_G) + (1 - \pi_R - \gamma)u(w_B)] \quad (3)
 \end{aligned}$$

where u is assumed to be twice differentiable, strictly increasing and strictly concave.¹⁰ The coefficient α measures aversion to the ambiguity inherent in the fact that true probabilities of states G, B are not known. Maximal aversion to ambiguity occurs at $\alpha = 1$; maximal loving of ambiguity occurs at $\alpha = 0$.¹¹ When $\alpha = 1/2$, the agent behaves like an expected utility maximizer with beliefs $(\pi_R, (1 - \pi_R)/2, (1 - \pi_R)/2)$, and so appears neutral with respect to ambiguity.¹²

Because we want to focus on ambiguity aversion, we assume $\alpha > 1/2$. Let $w = (w_R, w_G, w_B)$ be the optimal choice when prices are p . We begin by analyzing choice of wealth in the ambiguous states. If $w_G > w_B$, then the minimum in the formula (3) for utility occurs when $\beta = 0$ and the maximum occurs when $\gamma = 1 - \pi_R$, so utility is

$$U(w) = \pi_R u(w_R) + \alpha(1 - \pi_R)u(w_B) + (1 - \alpha)(1 - \pi_R)u(w_G)$$

¹⁰More generally, it might be assumed that the minimum and maximum are taken over sets of probabilities smaller than the entire interval $[0, (1 - \pi_R)]$.

¹¹When $\alpha = 1$ these preferences reduce to the “maxmin” preferences of Gilboa & Schmeidler (1989).

¹²This is a special implication of the fact that there are only three states, and R has known probability.

Hence the first order conditions for optimality are

$$\frac{\pi_R u'(w_R)}{p_R} = \frac{\alpha(1 - \pi_R)u'(w_G)}{p_G} = \frac{(1 - \alpha)(1 - \pi_R)u'(w_B)}{p_B} \quad (4)$$

Solving the last equality for the price ratio p_G/p_B and keeping in mind that $w_G > w_B$ and that u' is strictly decreasing, we obtain

$$\frac{p_G}{p_B} = \left[\frac{(1 - \alpha)}{\alpha} \right] \left[\frac{u'(w_G)}{u'(w_B)} \right] < \frac{(1 - \alpha)}{\alpha}$$

Conversely, if $w_B > w_G$ then the minimum in the expression for utility occurs when $\beta = 1 - \pi_R$ and the maximum occurs when $\gamma = 0$, so utility is

$$U(w) = \pi_R u(w_R) + \alpha(1 - \pi_R)u(w_G) + (1 - \alpha)(1 - \pi_R)u(w_B)$$

Hence the first order conditions for optimality are

$$\frac{\pi_R u'(w_R)}{p_R} = \frac{\alpha(1 - \pi_R)u'(w_B)}{p_B} = \frac{(1 - \alpha)(1 - \pi_G)u'(w_G)}{p_G} \quad (5)$$

and so we obtain

$$\frac{p_G}{p_B} = \left[\frac{\alpha}{(1 - \alpha)} \right] \left[\frac{u'(w_G)}{u'(w_B)} \right] > \frac{\alpha}{(1 - \alpha)}$$

Putting these together, we conclude that

$$\begin{aligned} w_G > w_B &\iff \frac{p_G}{p_B} < \frac{(1 - \alpha)}{\alpha} \\ w_G < w_B &\iff \frac{p_G}{p_B} > \frac{\alpha}{(1 - \alpha)} \\ w_G = w_B &\iff \frac{(1 - \alpha)}{\alpha} \leq \frac{p_G}{p_B} \leq \frac{\alpha}{(1 - \alpha)} \end{aligned} \quad (6)$$

Because $\alpha > 1/2$, the set of prices p where

$$\frac{(1 - \alpha)}{\alpha} < \frac{p_G}{p_B} < \frac{\alpha}{(1 - \alpha)} \quad (7)$$

is a non-empty *open* set. For such prices p , an ambiguity-averse agent (with α -maxmin utility) will insist on holding an *unambiguous portfolio* — that

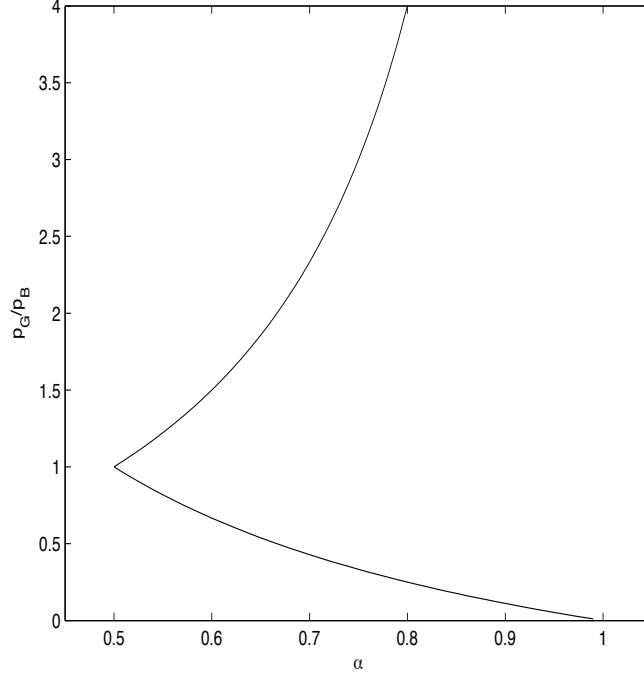


Figure 1: For $(\alpha, p_G/p_B)$ between the curves, the agent chooses an unambiguous portfolio

is, a portfolio with $w_B = w_G$; see Figure 1. If $\alpha = 1$ *all* prices satisfy the inequalities (7), so an agent who is maximally ambiguity averse will insist on holding an unambiguous portfolio for *all* prices.

We now turn to choice of wealth in the risky state. In the range of prices (7) the agent chooses $w_G = w_B$ so the first-order conditions are:

$$\frac{\pi_R u'(w_R)}{\pi_R} = \frac{(1 - \pi_R) u'(w_G)}{p_G + p_B} = \frac{(1 - \pi_R) u'(w_B)}{p_G + p_B} \quad (8)$$

(Because we shall make no use of them, we leave the derivation of the first order conditions outside this range of prices to the interested reader.)

2.3 Equilibrium Implications

The implications derived above for individual choice have immediate implications for equilibrium choices and hence for equilibrium prices.

Suppose first that all agents maximize expected utility with respect to a common prior $\pi = (\pi_R, \pi_G, \pi_B)$. At equilibrium, all agents face the same prices and individual choices w^n sum to the social endowment $W = \sum e^n$, so (2) implies that

$$W_\sigma > W_\nu \implies \frac{p_\sigma}{\pi_\sigma} < \frac{p_\nu}{\pi_\nu} \text{ and } w_\sigma^n > w_\nu^n \quad (9)$$

Suppose next that all agents are equally ambiguity averse (i.e., $\alpha^n = \alpha$ for each n) and $W_G \neq W_B$. If all agents are maximally ambiguity averse ($\alpha = 1$), then there cannot exist an equilibrium with positive prices, because maximally ambiguity averse agents always refuse to be exposed to ambiguity. If all agents are equally, but not maximally, ambiguity averse ($0.5 < \alpha < 1$) then it is easily seen that there is a unique equilibrium, having the property that prices and choices are exactly as they would be if all agents maximized expected utility with respect to the common prior

$$\hat{\pi}_\alpha = \left(\pi_R, (1 - \alpha)(1 - \pi_R), \alpha(1 - \pi_R) \right)$$

If $W_G < W_B$ we obtain the same conclusion except that the imputed prior is

$$\tilde{\pi}_\alpha = \left(\pi_R, \alpha(1 - \pi_R), (1 - \alpha)(1 - \pi_R) \right)$$

Suppose, finally, that attitudes toward ambiguity are heterogeneous across the population. In this situation, the description of equilibrium is much more complicated. To illustrate the point, suppose there are two types of agents: agents $\ell \in L$ (agents of *Type I*) maximize expected utility with respect to the common prior $\pi = (\pi_R, \pi_G, \pi_B)$, while agents $m \in M$ (agents of *Type II*) are maximally ambiguity averse.¹³ (Because there are only two types of agents, $\#L + \#M = N$.) Assume that $W_G \neq W_B$.

¹³The qualitative conclusions would be the same if we assumed that Type I agents were ambiguity neutral in the sense of maximizing .5-maxmin utility.

The discussion above yields implications for the distribution of equilibrium wealth in the ambiguous states, most easily seen by considering the distribution, across agents, of shares $w_G/(w_G + w_B)$ of each agent's wealth in state G as a fraction of that agent's total wealth in the ambiguous states.

(i) Because $\alpha = 1$, (6) implies that all Type II agents choose equal wealth in the ambiguous states: $w_G^m = w_B^m = w_a^m$. Hence $w_G^m/(w_G^m + w_B^m) = 1/2$ for all agents of Type II, and the distribution of wealth shares should have a mode at $1/2$.

(ii) Write

$$W_a^{II} = \sum_{m \in M} w_a^m$$

for the total wealth held in each of the ambiguous states by agents of Type II. Because markets clear in equilibrium, agents of Type I must hold, in aggregate, the remaining wealth:

$$\begin{aligned} W_G^I &= \sum_{\ell \in L} w_G^\ell = W_G - W_a^{II} \\ W_B^I &= \sum_{\ell \in L} w_B^\ell = W_B - W_a^{II} \end{aligned}$$

The weighted average of the wealth shares $w_G^\ell/(w_G^\ell + w_B^\ell)$ must equal the average net wealth shares, so assuming that $W_a^{II} > 0$, we have:

$$\frac{W_G^I}{W_G^I + W_B^I} = \frac{W_G - W_a^{II}}{(W_G - W_a^{II}) + (W_B - W_a^{II})} > \frac{W_G}{W_G + W_B}$$

Thus, the distribution of wealth shares for Type I agents should be skewed to the right of the distribution of wealth shares that would be expected absent ambiguity or ambiguity aversion.

These implications for choices have indirect implications for prices:

(iii) In view of (1), choices of Type I agents are sensitive to the entire vector p of state prices; in view of (4), (5) and (8), choices of Type II agents are

sensitive only to p_R and p_G+p_B . Put differently: all agents are marginal with respect to the determination of the price ratio $p_R/(p_G + p_B)$ but only Type I agents are marginal with respect to the determination of the price ratio p_G/p_B .

- (iv) As noted above, ambiguity averse (Type II) agents choose to hold equal wealth w_a^m in the ambiguous states G, B and choose wealth w_R^m in the risky state R according to the budget constraint and the first-order condition (8). Note that each agent's state-dependent wealth need *not* be ranked opposite to the ranking of state/price probabilities, and hence that aggregate state-dependent wealth held by Type II agents also need *not* be ranked opposite to the ranking of state price/probabilities. In view of our earlier discussion, aggregate state-dependent wealth held by agents of Type I *must* be ranked opposite to the ranking of state price/probabilities. Because agents of Type I hold the difference between aggregate state-dependent wealth and aggregate state-dependent wealth held by agents of Type II, it follows that state price/probabilities need *not* be ranked opposite to the ranking of aggregate state-dependent wealth.

What rankings are possible? If $W_G > W_B$ and agents of Type II choose equal wealth in the ambiguous states G, B then $W_G^I > W_B^I$. In view of (1), the wealth choices of Type I agents should be ranked opposite to state price/probability ratios; because these choices sum to W_G^I and W_B^I , it follows that state price/probability ratios should be ranked opposite to social wealth: $p_G/\pi_G < p_B/\pi_B$. As the reader can easily see, no matter how aggregate wealth in the risky state W_R is ranked with respect to aggregate wealth in the ambiguous states W_G, W_B , any ranking of the state price/probability ratio for the risky state p_R/π_R with respect to the state price/probability ratios for the ambiguous states $p_G/\pi_G, p_B/\pi_B$ is theoretically possible. However, not all rankings are equally likely. For example, if we consider environments with a single agent of each type, we find that aggregate wealth rankings $W_G > W_R > W_B$ are less likely to lead to rankings of state price/probabilities that are different from the predictions when all

agents maximize expected utility than are aggregate wealth rankings $W_G > W_B > W_R$.

Finally, we note an implication for representative agent pricing. If all agents are ambiguity neutral then the ranking of state price/probabilities is opposite to the ranking of aggregate wealth, and prices can be rationalized by the preferences of a representative agent who maximizes expected utility with respect to the common prior.¹⁴ If some agents are ambiguity averse and the ranking of state price/probabilities is not opposite to the ranking of aggregate wealth, prices *cannot* be rationalized by the preferences of a representative agent who maximizes expected utility with respect to the common prior.

¹⁴See Constantinides (1982) for example.

3 Experimental Design

The following is a brief description of our experimental design and of the parameters for each of the ten experimental sessions.

Each experimental session consisted of a sequence of eight *trading periods*, of fixed and announced length. At the beginning of each trading period, subjects were endowed with securities and cash.¹⁵ During each trading period, markets were open and subjects were free to trade securities, using cash as the means of exchange. At the end of the trading period, markets closed, the state of the world was revealed, and security dividends were paid. Dividends of end-of-period holdings of securities and cash constituted a subject's period earnings, but actual payments were only made at the end of the experiment.¹⁶ (Thus earnings in each period did not affect endowments in future periods.) At the end of the experimental session, the cumulated period earnings were paid out to the subject, together with a sign-up reward.

Two kinds of securities, bonds and stocks, were traded. Bonds paid a fixed dividend of \$0.50. Stocks paid a random dividend, depending on the state of the world: the Red (respectively Green, Blue) security paid \$0.50 if the state was revealed to be Red (respectively Green, Blue) and nothing otherwise. Subjects were allowed to short-sell stocks and bonds, as long as they did not take positions that could result in losses of more than \$2.00.¹⁷

¹⁵In some sessions, cash and security payoffs were denominated in US dollars; in other sessions, cash and security payoffs were denominated in a fictitious currency called *francs*; at the end of the session, francs were converted to dollars at a pre-announced rate. The results do not appear to depend on the denomination of payoffs.

¹⁶In some sessions, some subjects were given a loan of cash which they were required to repay from end-of-period proceeds; in other sessions, subjects received a negative endowment of bonds — a loan, in a different guise. Here we report loans as negative endowments of bonds.

¹⁷In the early sessions, we imposed this limit *ex post*, by barring subjects with more than \$2.00 losses from trading in future periods. In later sessions we employed software that checked pending orders against a *bankruptcy rule*: wealth was computed in all possible states, assuming that all orders within 20% of the best bid or ask were executed; if losses were larger than \$2.00, the pending order was rejected.

Trading took place over an electronic market organized as a continuous open-book double auction in which infra-marginal orders remained displayed until executed or canceled.¹⁸

The state of the world was determined by a draw from an urn. Initially, the urn contained 18 balls, of which 6 were Red and the others were either Green or Blue. In some sessions, subjects were told the entire composition of the urn, so the environment was one of pure risk; in other experiments subjects were told only the total number of balls and the number of Red balls, so the environment involved both risk and ambiguity. Balls were drawn without replacement, so both the total number of balls in the urn and the number of balls of each color (and hence the proportion of balls of each color) changed during the course of the experimental sessions.

Sessions typically lasted 2.5 hours and began with two practice periods.¹⁹ Subject earnings ranged from \$0 to \$125, with an average of approximately \$50.

We conducted experimental sessions distinguished by the endowment distribution, the urn composition, and the ambiguity/risk environment. As discussed in Subsection 2.3, endowment distributions were chosen so that reversals in the ranking of state price/probabilities were more or less likely: PRR = possible (more likely) rank reversals; NRR = no (less likely) rank reversals. For each of the two choices of endowment distributions we conducted experimental sessions with three different urn compositions (A, B,

¹⁸Three different interfaces were used: (i) *Marketscape* (developed in Charles Plott's lab), in which quantities and prices had to be entered manually; (ii) *eTradeLab* (developed by Tihomir Asparouhov) in which market orders (orders that executed immediately at the best available price) could be entered by clicking only; (iii) *jMarkets* (developed at Caltech, and available as open source software at <http://jmarkets.ssel.caltech.edu>) in which all orders were submitted by point-and-click. The results do not appear to depend on the interface used.

¹⁹In some experiments subjects were paid in practice periods and in some experiments subjects were not paid in practice periods, but in neither case are the results from practice periods recorded in the data. The results do not appear to depend on payments in practice periods.

C). Finally, for four of the sessions (corresponding to four vectors of endowment distributions/urn compositions), we repeated the session with different subjects but with the *same* endowment distributions, the same urn compositions and the *same* sequence of draws from the urn — but we *announced* the true composition of the urn. Thus, we created four sets of paired sessions in which it is possible to compare outcomes in environments with ambiguity and environments with pure risk. For convenience, we identify each of the ten experiments by the endowment distribution, urn composition and ambiguity/risk treatment; eg., (NRR, B, Risk). For the various sessions, Table 1 shows the security endowments for each subject type, Table 2 shows the corresponding wealth distributions, Table 3 shows the number of subjects of each type, and Table 4 shows the fraction of aggregate wealth in each of the ambiguous states, computed from the endowments for each subject type and the number of subjects of each type. (The numbers shown are approximate, because they differ slightly according to the precise number of subjects of each type.) Finally, Table 5 shows the urn composition.

Subjects were told their own endowments but *not* the endowments of others or aggregate endowments. In particular, subjects had no way of knowing the ranking of wealth, and so could not distinguish between the treatments NRR and PRR.

Details of the last experimental session — the pure-risk replication of Treatment (NRR,B) — can be viewed on the experimental web site:

<http://clef.caltech.edu/exp/amb/start.htm>

Table 1: Security Endowments

Experiment Type	Subject Type	Endowments				
		R	G	B	Notes	Cash
NRR	i	0	28	0	0	\$2.00
	ii	20	12	0	-5	\$1.00
NRR	i	3	12	0	0	\$1.20
	ii	0	4	9	-4	\$2.00

Table 2: Initial Wealth

Experiment Type	Subject Type	Wealth		
		R	G	B
NRR	i	\$2.00	\$16.00	\$2.00
	ii	\$8.50	\$4.50	-\$1.50
PRR	i	\$2.70	\$7.20	\$1.20
	ii	\$0.00	\$2.00	\$6.50

Table 3: Number of Subjects of Each Type

Experiment Type	Experiment	Risk	Ambiguity
NRR	A	(15,14)	(15,14)
	B	(15,14)	(15,14)
	C	-	(13,13)
PRR	A	(15,14)	(13,13)
	B	(12,12)	(12,12)
	C	-	(15,14)

Table 4: Aggregate Wealth Distribution in the Ambiguous States (Approximate)

Experiment Type	$W_G/(W_G + W_B)$	$W_B/(W_G + W_B)$
NRR	0.98	0.02
PRR	0.63	0.37

Table 5: Initial Composition of Urn

Experiment Type	Experiment	Urn		
		R	G	B
NRR	A	9	3	6
	B	6	3	9
	C	9	3	6
PRR	A	6	6	6
	B	6	6	6
	C	6	6	6

4 Empirical Findings

In this Section we discuss the experimental data, first with regard to the cross-sections of security holdings and then with regard to the state securities' price/probability ratios. In the last Subsection, we discuss possible correlation between risk aversion and ambiguity aversion.

4.1 End-of-period Wealth

Perhaps the clearest evidence of the existence and effect of ambiguity aversion is to be found in the cross-sectional distribution of end-of-period wealth. Perhaps the clearest and most striking way to see the effect is to compare end-of-period wealth in the four paired risk/ambiguity treatments: (NRR, A, Risk) and (NRR, A, Ambiguity); (NRR, B, Risk) and (NRR, B, Ambiguity); (PRR, A, Risk) and (PRR, A, Ambiguity); (PRR, B, Risk) and (NRR, B, Ambiguity).

These comparisons are presented as histograms in Figures 2 and 3. In each Figure, the top two panels provide the results for the Risk treatments (left: configuration A; right: configuration B) and the lower panels provide the results for the corresponding Ambiguity treatments. In each of the four Risk treatments, the observed distribution of $w_G/(w_G + w_B)$ (individual wealth in the Green state as a proportion of individual wealth in the two ambiguous states) is nearly uni-modal, and very consistent with the market wealth ratios $W_G/(W_G + W_B)$ (which are approximately .98 in the NRR treatments and .63 in the PRR treatments). However, in the four Ambiguity treatments the modes have shifted to .50, apparently reflecting choices of ambiguity-averse subjects. Moreover, the distributions have significantly bigger right tails, reflecting the compensating choices of ambiguity-tolerant subjects.

For the profiles (NRR, C, Ambiguity) and (PRR, C, Ambiguity) we have no corresponding paired Risk treatments. However, as Figure 4 shows, the

distributions of end-of-period wealth are entirely consistent with theory, with the mode at .50 and heavy right tails.

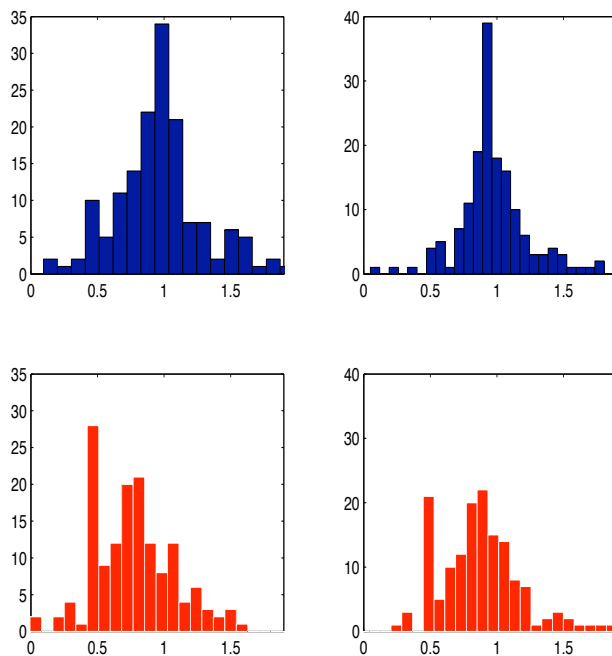


Figure 2: Histograms of final wealth in state G as a proportion of final wealth in states G and B , NRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.

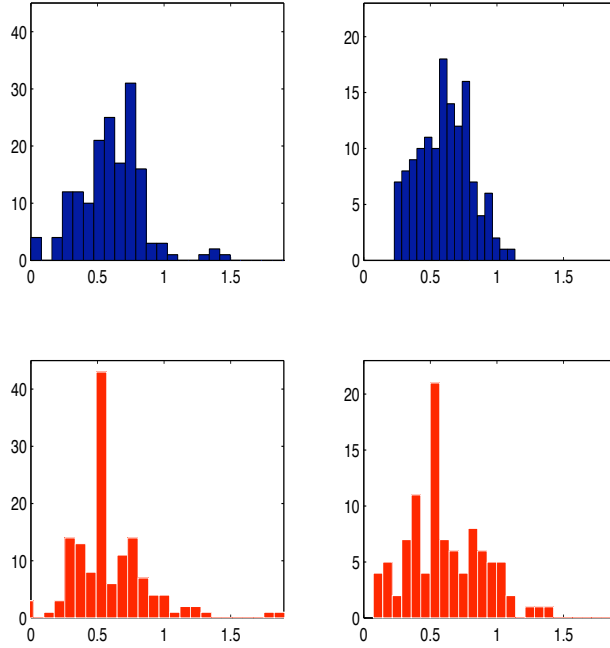


Figure 3: Histograms of final wealth in state G as a proportion of final wealth in states G and B , PRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.

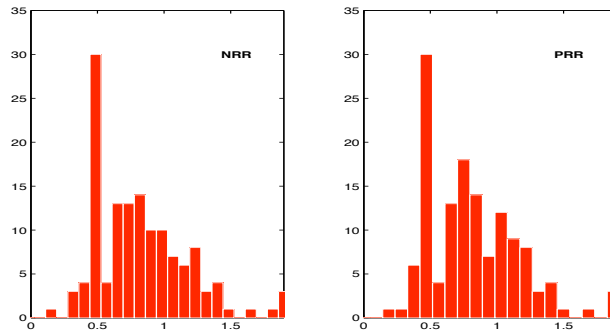


Figure 4: Histograms of final wealth in state G as a proportion of final wealth in states G and B ; (left: NRR treatment, right: PRR treatment).

4.2 State Price/Probability Ratios

By definition, state price/probability ratios are the ratios of state prices to state probabilities. In the Risk treatments, the probabilities π_R, π_G, π_B are known, so state price/probability ratios are easily computed. However, in the Ambiguity treatments, only π_R is known, so it is not obvious which state probabilities to use in computing state price/probability ratios for the ambiguous states G, B . Here we follow the simplest approach and use uniform priors over the ambiguous states for the initial draw, updated by Bayes' Rule for subsequent draws. (Other choices are certainly possible, but would not yield uniformly better results; see Section 5 for discussion.)

We display pricing results in the form of empirical cumulative distribution functions (ECDFs), for several reasons. The first, and perhaps most important, reason is that ECDFs provide unbiased estimates, unaffected by time series considerations such as autocorrelation and conditional heteroscedasticity, of the probability that a state price/probability ratio exceeds any given level. That is, ECDFs provide unbiased answers to questions of the type

$$\text{Is } \text{Prob}(p_R/\pi_R > 1) > \text{Prob}(p_B/\pi_B > 1) ?$$

Because the Glivenko-Cantelli theorem implies that ECDFs converge uniformly to the true underlying distribution, focusing on ECDFs means that questions concerning first-order stochastic dominance such as

$$\text{Is } \text{Prob}(p_R/\pi_R > a) > \text{Prob}(p_B/\pi_B > a) \text{ for every } a ?$$

are meaningful. The second reason is that we have no direct knowledge of subjects' actual attitudes towards risk and ambiguity, and so focus on ordinal comparisons. Finally, because markets go through lengthy adjustments — even in situations as simple as the present ones — many (perhaps most) transactions take place before markets “settle.” (In fact, in some experiments, it is not clear that markets ever settled.)

As above, we focus on the paired Risk/Ambiguity treatments, as they allow us to make the sharpest comparisons.

First, consider the NRR treatment. By construction, $W_G > W_R > W_B$, so in the Risk treatment theory predicts $p_B/\pi_B > p_R/\pi_R > p_G/\pi_G$. Moreover, as the discussion in Subsection 2.3 suggests, we predict that the same ordering should be most likely in the Ambiguity treatment as well. As Figure 5 shows, this is what we see in the data. (The top panels of Figure 5 display ECDFs for the NRR Risk treatments and the bottom panels display ECDFs for the corresponding NRR Ambiguity experiments.) In both cases, the state price/probability ratio for B stochastically dominates the state price/probability ratio for R and the state price/probability ratio for R stochastically dominates the state price/probability ratio for G .

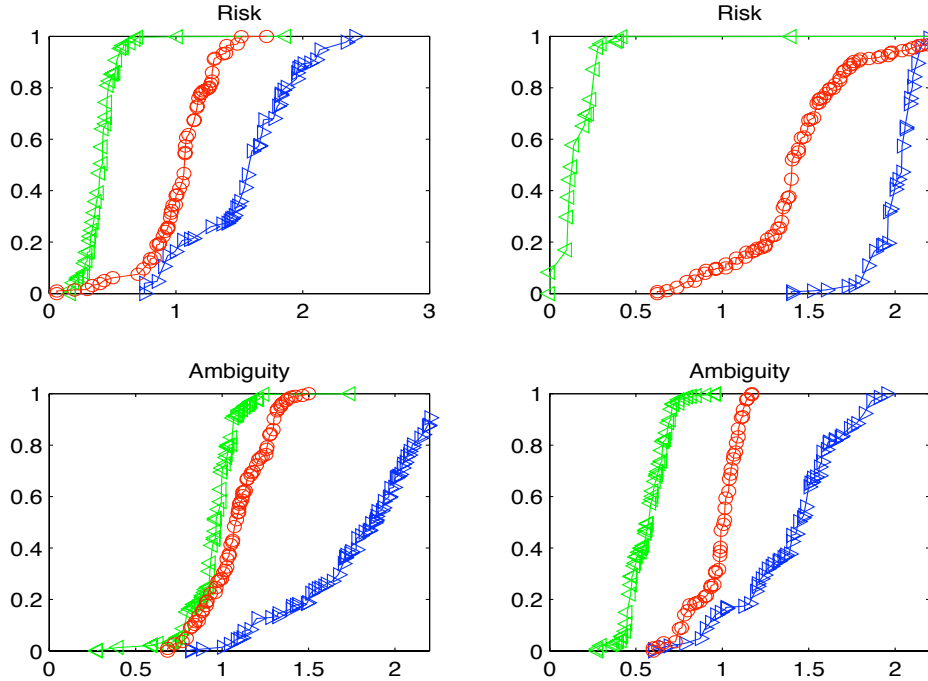


Figure 5: Empirical Distribution Functions (ECDFs) of state price/probability ratios, NRR treatment. Top panels: pure risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment. Distribution with (green) arrows pointing to the left is for state G ; distribution with (blue) arrows pointing to the right is for state B ; distribution with (red) circles is for state R .

Now consider the PRR treatment. By construction, $W_G > W_B > W_R$, so in the Risk treatment theory predicts $p_R/\pi_R > p_B/\pi_B > p_G/\pi_G$. However, the discussion in Subsection 2.3 suggests that we may see a rank reversal, leading to the ordering $p_B/\pi_B > p_R/\pi_R > p_G/\pi_G$. As the left panels of Fig-

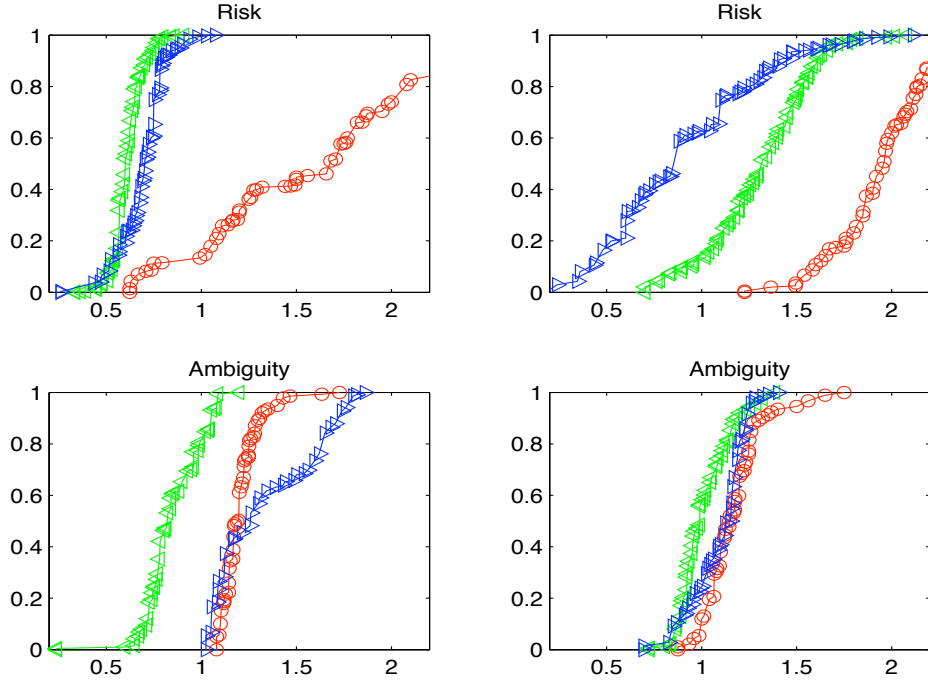


Figure 6: Empirical Distribution Functions (ECDFs) of state price/probability ratios, PRR treatment. Top panels: pure risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment. Distribution with (green) arrows pointing to the left is for state G ; distribution with (blue) arrows pointing to the right is for state B ; distribution with (red) circles is for state R .

ure 6 show, for the A version this is pretty much what we see in the data. In the Risk treatment, the state price/probability ratio for R dominates the ratio for B , which in turn dominates the ratio for G ; in the Ambiguity treatment the state price/probability ratios for B and R dominate the state/price probability ratio for G and the ECDF for B is to the right of the ECDF for R most (although not all) of the time. In the (Risk, B) version, we see

anomalous rankings: the ECDF for B is to the left of the ECDF for G much of the time. Such violations have been observed before (Bossaerts & Plott, 2004) when, as happened here, an unusual sequence of draws occurred. In this case, B was drawn four times in six periods and G was never drawn at all. As a result, in later periods, subjects seemed to believe G was much more more likely to be drawn and B was much less likely to be drawn, driving p_G up and p_B down. In the corresponding Ambiguity treatment, the ECDF for B is shifted upward and very close to the ECDF for R and the ECDF of G is to the left of the ECDF for B ; the appreciation of p_B is consonant with what we would expect in the presence of ambiguity averse subjects.²⁰

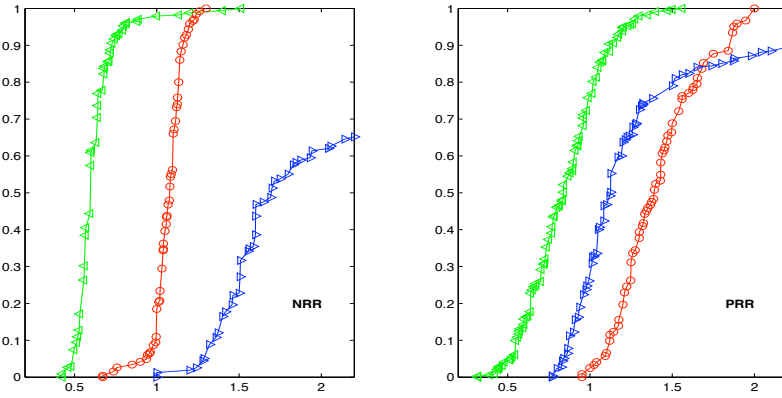


Figure 7: Empirical Distribution Functions (ECDFs) of state price/probability ratios. Left = (NRR, C, Ambiguity); Right = (PRR, C, Ambiguity). Distribution with (green) arrows pointing to the left is for state G ; distribution with (blue) arrows pointing to the right is for state B ; distribution with (red) circles is for state R .

The data for the final two sessions are shown in Figure 7. In the left panel, for which the configuration is (NRR, C, Ambiguity), we expect, and see, the rankings $p_B/\pi_B > p_R/\pi_R > p_G/\pi_G$, just as if probabilities were known. In the right panel, for which the configuration is (PRR, C, Ambiguity), the predicted rankings would be $p_R/\pi_R > p_B/\pi_B > p_G/\pi_G$ if probabilities were known; in the actual data the rankings appear anomalous.

²⁰This kind of problem could be avoided by exercising some control over the sequence of draws — but then the draws would no longer be random.

In sum, the pricing effects from the introduction of ambiguity are consistent with the theoretical analysis of Section 2: rank changes in state price/probability ratios are observed only in the PRR treatment, when the security in shortest supply pays off in a risky, rather than ambiguous, state of the world.

4.3 Risk Aversion and Ambiguity Aversion

In principle, there seems to be no reason why risk aversion (in our framework, concavity of the felicity function u) and ambiguity aversion (in our framework, the coefficient α) need be coordinated. However, our experimental data suggests that they may in fact be positively correlated.

We compare the range of end-of-period wealth across all states — which is a measure of risk tolerance — with the range of end-of-period wealth across the ambiguous states — which is a measure of ambiguity tolerance. Figure 8 displays the results for all subjects and all periods in all the sessions what involved ambiguity. We see a significant positive correlation between risk tolerance (a wide range of end-of-period wealth in all states) and ambiguity tolerance (a wide range of end-of-period wealth in the ambiguous states).²¹

A significant positive correlation between ambiguity and risk aversion would have serious implications for asset pricing. It would suggest, for instance, a novel explanation of the value effect — the observation that securities in companies with high book-to-market values earn higher returns (equivalently, carry a higher risk premium) than securities in companies with low book-to-market values. Low book-to-market value suggests growth potential, and hence greater ambiguity about future performance. Hence securities with low book-to-market values should be held mostly by ambiguity tolerant agents, while securities with high book-to-market values should be held by a broader mix of investors. If ambiguity tolerant agents are also more risk tolerant, then they require a lower risk premium, so the return on

²¹Our findings are consistent with at least one study in neuroscience (Hsu *et al.*, 2005).

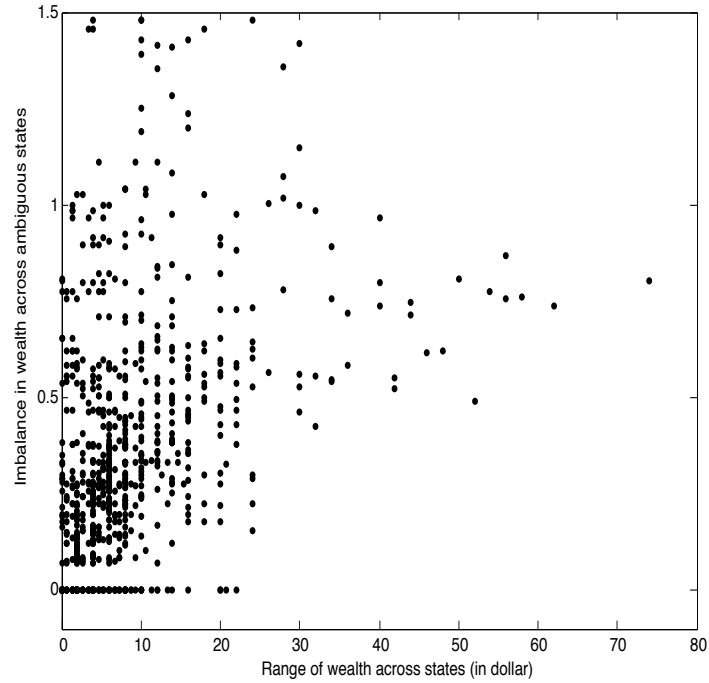


Figure 8: Plot of difference from 0.5 of wealth allocated to state g as a proportion of final wealth in states g and b (the ambiguous states), against range of wealth allocated across all states; all periods in all experiments with ambiguous states.

securities with low book-to-market values (growth stocks) should be lower than the return on securities with high book-to-market values (value stocks).

Correlation between ambiguity and risk aversion might also be relevant for regulation (Easley & O'Hara, 2005).

5 Discussion

There are two issues that deserve more discussion. The first concerns alternative explanations for choice and price patterns, and in particular, whether (heterogeneous) ambiguity aversion is really required to explain the experimental data. The second concerns our choice to use a “uniform” probability on ambiguous states (more precisely, beginning with uniform priors on the ambiguous states and then using Bayesian updating) in calculating state price/probabilities.

As we have already noted, if risk attitudes are heterogeneous across the population but beliefs and ambiguity attitudes are homogeneous (and in particular if all agents are ambiguity neutral), then pricing can be rationalized by the existence of a representative agent who shares the common beliefs and whose utility function is state-independent. The pricing data in the NRR treatments do not contradict the existence of such a representative agent. Matters are less clear for the PRR treatments because it is not obvious what probabilities to use in computing state/price probability ratios. As discussed, when we assume a “uniform” prior we sometimes observe ECDFs for state price/probabilities that are *inconsistent* with the existence of a representative agent. For each given experimental session, it does seem possible to find some priors for which the ECDFs are consistent with the existence of a representative agent. (Indeed, it is not clear that any data from a single experimental session could ever be inconsistent with the existence of a representative agent with an appropriately chosen prior probability.) But it does not seem possible to find any set of priors that yields ECDFs for state price/probabilities that are consistent with the existence of a representative agent in *all* experimental sessions.

Leaving aside the obvious question of why we might observe *different* commonly held priors in experimental sessions in which the information given to subjects was *the same*, it is the patterns of final holdings that are most difficult or impossible to explain on the basis of heterogeneity in risk attitudes. As we have discussed, in each experimental session where probabilities of

the G, B states were not known, there was a sizable group of subjects who chose not to be exposed to ambiguity; i.e., subjects who chose $w_G = w_B$. This is the behavior that would be expected of (extremely) ambiguity averse agents. Heterogeneity in risk attitudes by itself cannot explain that, as it would imply that all agents' wealth should be comonotonic with the social wealth, a prediction that is hard to square the presence of so many subjects who choose equal wealth in the ambiguous states — especially in the NRR treatments, in which the social endowments W_G, W_B and prices p_G, p_B are far apart.

Alternatively, one might imagine that heterogeneity *of prior beliefs* (which seems especially natural in the context of unknown probabilities), might be enough to explain our observations. If *all* agents are good Savage Bayesians and so in particular maximize expected utility with respect to some prior, but hold different priors, should we expect to see prices and holdings like the ones we see in our experimental data? As the discussion above makes clear, the pricing data will be consistent with any model which admits a representative agent with *some* priors, and as we have noted earlier, it seems that we can always find such a representative agent and such priors. However, to believe that this is a satisfactory model, we would have to be prepared to believe that in the NRR treatments the subjects' priors are such that the representative agent has “uniform” priors, while in the PRR treatments the subjects' priors are such that the representative agent has priors different enough from “uniform” priors to yield the observed ranking of ECDFs — despite the fact that subjects *do not know the treatment*.

Heterogeneity of prior beliefs is even less successful in explaining the experimental data on choice (whether or not a representative agent exists). As we observed in Section 2, an agent who maximizes subjective expected utility will choose equal wealth in the ambiguous states ($w_G = w_B$) only when the subjective state price/probability ratios of the two states are equal. It would seem to be a remarkable coincidence to observe in *every* experimental session a large group of subjects whose priors imply equal subjective state price/probability ratios for the ambiguous states and another large group of subjects whose priors imply quite different subjective state price/probability

ratios for the ambiguous states. However, because agents who are ambiguity averse will choose equal wealth in the ambiguous states for an *open* set of prices, this is exactly what we would expect to see in a world in which a significant fraction of agents are ambiguity averse and a significant fraction are ambiguity neutral.

6 Conclusion

The most important findings of this paper are that ambiguity aversion can be observed in competitive markets and that ambiguity aversion matters for portfolio choices and for prices. The predictions for portfolio choices seem quite robust and well-supported by the experimental data; the predictions for prices are less robust. This is a somewhat surprising state of affairs: much of asset pricing theory claims to make sharp predictions about prices but much less sharp predictions about portfolio choices. For a related discussion, see Bossaerts, Plott & Zame (2007).

Our theoretical and experimental findings contradict two apparently widespread and often-asserted beliefs. The first is that that prices reflect an average of the beliefs of *all* agents.²² In our setting, agents who are sufficiently ambiguity averse choose not to be exposed to ambiguity, so their beliefs about ambiguous states are *not* reflected in prices. The second is that infra-marginal agents have no effect on prices. In our setting, the ambiguity averse infra-marginal agents do not have a direct effect on the prices of ambiguous securities, but they do affect the amount of risk held by the ambiguity neutral marginal agents and hence have an indirect effect on prices.

²²See Hirshleifer (2001) for instance.

References

- Bossaerts, P. and C. Plott (2004): “Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets.” *Review of Finance* 8, 135-169.
- Bossaerts, P., C. Plott and W. Zame (2007): “Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiment,” *Econometrica* 75, 993-1038.
- Cagetti, M., Hansen, L., T. Sargent and N. Williams (2002): “Robustness and Pricing with Uncertain Growth.” *Review of Financial Studies* 15, 363-404.
- Calvet, L. and Grandmont, J.M., and I. Lemaire (2002): “Aggregation of Heterogeneous Beliefs and Asset Pricing in Complete Financial Markets,” Harvard University, working paper.
- Chapman, D. and V. Polkovnichenko (2005): “Heterogeneity in Preferences and Asset Market Outcomes,” Carlson School of Management, University of Minnesota, working paper.
- Constantinides, G., “Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation,” *Journal of Business* 55, 253-267.
- Dana, R.A. (2004): “Ambiguity, Uncertainty Aversion and Equilibrium Welfare.” *Economic Theory* 23, 569-587.
- Easley, D. and M. O’Hara (2005): “Regulation and Return: The Role of Ambiguity,” Johnson GSM, Cornell University, working paper.
- Ellsberg, D. (1961): “Risk, Ambiguity and the Savage Axioms.” *Quarterly Journal of Economics* 75, 643-669.
- Epstein, L. and T. Wang (1994): “Intertemporal Asset Pricing Under Knightian Uncertainty.” *Econometrica* 62, 283-322.

- Epstein, L. and J. Miao (2003): "A Two-Person Dynamic Equilibrium Under Ambiguity." *Journal of Economic Dynamics and Control* 27, 1253-1288.
- Fama, E. and K. French (1992): "The Cross-Section of Expected Stock Returns." *Journal of Finance* 47, 427-465.
- Fama, E. and K. French (1992): "Value versus Growth: The International Evidence." *Journal of Finance* 53, 1975-1999.
- Fehr, E. and J.R. Tyran (2005): "Individual Irrationality and Aggregate Outcomes," *Journal of Economic Perspectives* 19, 43-66.
- Ghirardato, P., F. Maccheroni and M. Marinacci (2004): "Differentiating Ambiguity and Ambiguity Attitude." *Journal of Economic Theory* 118, 133-173.
- Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with a Non-Unique Prior." *Journal of Mathematical Economics* 18, 141-153.
- Gneezy, U., A. Kapteyn and J. Potters (2003): "Evaluation Periods and Asset Prices in a Market Experiment." *Journal of Finance* 58, 821-837.
- Hirshleifer, D. (2001): "Investor Psychology and Asset Pricing," *Journal of Finance* 56, 1533-1597.
- Hsu, M., M. Bhatt, R. Adolphs, D. Tranel and C. Camerer (2005): "Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making," *Science* 310, 1680-1684.
- Jouini, E. and C. Napp (2006): "Consensus Consumer and Intertemporal Asset Pricing with Heterogeneous Beliefs." *Review of Economic Studies*, forthcoming.
- Kluger, B. and S. Wyatt (2004): "Are Judgment Errors Reflected in Market Prices and Allocations? Experimental Evidence Based on the Monty Hall Problem." *Journal of Finance*, forthcoming.

- Knight, Frank (1939): *Risk, Uncertainty and Profit*, London: London School of Economics.
- Maenhout, P. (2000): “Robust Portfolio Rules and Asset Pricing.” INSEAD working paper.
- Savage, L.J. (1954): *The Foundations of Statistics*, J. Wiley and Sons, New York.
- Skiadas, C. (2005): “Dynamic Portfolio Choice and Risk Aversion,” in: *Handbook of Financial Engineering*, forthcoming.
- Trojani, F., M. Leippold and P. Vanini (2005): “Learning and Asset Prices under Ambiguous Information,” Sankt Gallen University working paper.
- Uppal, R. and T. Wang (2003): “Model Misspecification and Under Diversification.” *Journal of Finance*, forthcoming.