

Ambiguity in Asset Markets:
Theory and Experiment

Peter Bossaerts

Paolo Ghirardato

Serena Guarnaschelli

William R. Zame

No.27, September 2006*

www.carloalberto.org

*Revised in March 2009

Ambiguity in Asset Markets: Theory and Experiment¹

Peter Bossaerts² Paolo Ghirardato³ Serena Guarnaschelli⁴
William R. Zame⁵

March 2003; This Version March 2009⁶

¹An 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 the editor and the referee, from audiences at Alicante, Bocconi, CIRANO, Collegio Carlo Alberto, Columbia, East Anglia, Iowa, Irvine, Johns Hopkins, Kobe, Mannheim, Mainz, Northwestern, Nottingham, Paris 1, Salerno, Sankt Gallen, Southampton, 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, the Collegio Carlo Alberto or any other funding agency.

²California Institute of Technology & Ecole Polytechnique Fédérale Lausanne

³DSMA & Collegio Carlo Alberto, Università di Torino

⁴McKinsey, Inc.

⁵University of California, Los Angeles

⁶© 2009 by Peter Bossaerts, Paolo Ghirardato, Serena Guarnaschelli and William R. Zame.

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.

JEL Classification: C91, D53, D81, G11, G12

Keywords: Ambiguity, Experiments, Financial Markets, Heterogeneity

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 with respect to such prior.

Our experimental findings confirm the predictions of this theoretical analysis. We find that a significant fraction of agents are sufficiently ambiguity averse that they refuse to hold an ambiguous portfolio, that the degree of ambiguity aversion is heterogeneous across the population, and that rankings of state price/probability ratios can be anomalous, and are anomalous exactly in those configurations when theory suggests 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. (Such behavior violates the Savage (1954) independence axiom.)

We embed 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 love.

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 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.

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.

If supplies of ambiguous securities (Blue and Green, in our case) are unequal and the degree of ambiguity aversion is heterogeneous across the population, this choice behavior has an immediate implication for the cross-section of equilibrium portfolio holdings. Because sufficiently ambiguity averse agents will choose to hold an unambiguous portfolio, the remaining — 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.

In the same context, this choice behavior also has a more subtle implication for equilibrium pricing. If all agents maximize expected utility with respect to a common prior (but with possibly different risk attitudes), equilibrium state price/probability ratios will be ranked oppositely to aggregate wealth. However, if some agents are sufficiently ambiguity averse the situation may be quite different, even if their (generalized) “beliefs” coincide. As noted above, agents who are sufficiently averse to ambiguity will choose to hold an unambiguous portfolio. If the ambiguous securities are in unequal total supply, this means that the remaining — less ambiguity averse — agents must hold the imbalance of ambiguous securities. Because the relative prices of these ambiguous securities are determined by the marginal rates of substitution of these remaining agents, this effect tends to distort state price/probability ratios; if the distortion is sufficiently large, state price/probability ratios may not be ranked opposite to aggregate wealth.

This implication for rankings of state/price probabilities in turn has an implication for representative agent pricing. If all agents maximize expected utility with respect to a common prior (but with possibly different risk attitudes), equilibrium prices can always be rationalized by a representative agent who maximizes expected utility with respect to that common prior.⁵ However, if the distortion created by the presence of ambiguity averse agents is sufficiently large that state price/probability ratios are not be ranked opposite to aggregate wealth, equilibrium prices *cannot* be rationalized by a

⁵See Constantinides (1982), for example.

representative agent who maximizes expected utility with respect to any “obvious” prior — or even by a representative agent who is ambiguity averse and who has any “obvious” 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 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. (In some experiments, there is 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.

Our finding of a mode at equal portfolio weights under ambiguity is also interesting as it bears on a theoretical debate. The question is whether ambiguity aversion is to be modelled in a non-smooth fashion, as we do following Gilboa & Schmeidler (1989), or in a smooth fashion, as is done in the aptly called “smooth ambiguity” model of Klibanoff, Marinacci & Mukerji

⁶In the text we discuss in some detail what we mean by “obvious” priors and beliefs.

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

(2005), which is receiving increasing attention in the finance literature (see, e.g., Izhakian & Benninga (2008)). For an agent with smooth ambiguity preferences, equal portfolio weights are — as is the case with expected utility maximizing agents — a knife-edge choice. So we do not think that such agents significantly contribute to the mode at equal weighting observed in our experiments with ambiguity.⁸

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. Assuming, as seems natural, that value stocks are more like risky securities and growth stocks are more like ambiguous securities, 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 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.⁹

The approach here follows Bossaerts, Plott & Zame (2007), who study environments with pure risk (i.e., known probabilities). That paper documents 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 obtains (approximately) even though port-

⁸ A similar finding is reported, in a pure choice context, by Ahn, Choi, Gale & Kariv (2009).

⁹We thank Nick Barberis for this observation.

folio separation does not. (The market portfolio is the modal holding, but not at all the universal holding.) 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.

Two recent papers provide notable complements to our work. Easley & O'Hara (2009) 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 goes on to derive (theoretical) implications for regulation, assuming that risk aversion and ambiguity aversion are uncorrelated. However, these authors provide no experimental or historical data or empirical analysis to suggest that their assumptions about risk aversion and ambiguity aversion or their theoretical predictions are actually observed. Ahn, Choi, Gale & Kariv (2009) uses experimental data to estimate the extent of ambiguity aversion in a subject population. However, these authors work entirely in an individual choice environment, rather than in a market environment.

The above-mentioned papers belong to an emerging literature that studies the impact of non-expected utility preferences on prices and choices in competitive markets, either theoretically or experimentally. 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 (2006) study the effects of a particular class (rank-dependent expected utility) of non-expected utility 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 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 (2004), Skiadas (2008), Trojani, Leippold & Vanini (2007), seeks to explain the equity

premium puzzle (high average returns on equity and low average riskfree rate) by appealing to ambiguity (which they call “Knightian” or “model” uncertainty) on the basis of a model with an ambiguity averse representative agent. However, we show that ambiguity aversion does not aggregate across a heterogeneous population, so that prices may not be rationalizable by an ambiguity averse representative agent. Hence, our finding of 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 .¹⁰ 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 $p \cdot w^n \leq p \cdot e^n$

¹⁰Obviously 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 a Bernoulli utility function (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

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

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:

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

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 magnitude of prices.

2.2 Individual Choice: Ambiguity Aversion

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

To model preferences that are sensitive to ambiguity, we employ a generalization of Gilboa & Schmeidler (1989) called the α -*maxmin model*; see Ghirardato, Maccheroni & Marinacci (2004) for an axiomatization and more detailed discussion. We assume that there are a Bernoulli utility function u (assumed to be twice differentiable, strictly increasing and strictly concave), a convex set C of subjective priors on the state space $S = (R, G, B)$, and an $\alpha \in [0, 1]$ such that the utility of a wealth profile $w = (w_R, w_G, w_B)$ is assessed by

$$U(w) = \alpha \min_{\pi \in C} \left[\sum_{\sigma \in S} u(w_\sigma) \pi_\sigma \right] + (1 - \alpha) \max_{\pi \in C} \left[\sum_{\sigma \in S} u(w_\sigma) \pi_\sigma \right] \quad (3)$$

If C is a singleton, this reduces to subjective expected utility, so the extent to which C is not a singleton is a reflection of the perceived degree of ambiguity. The coefficient α measures the degree of aversion to this perceived ambiguity: maximal aversion to ambiguity occurs at $\alpha = 1$ (corresponding to the “maxmin” preferences of Gilboa & Schmeidler (1989)); maximal loving of ambiguity occurs at $\alpha = 0$.

In our setting, it is natural to assume that the event of a draw of a red ball from the urn is unambiguous with given probability π_R , so the set C reduces to an interval in the two dimensional unit simplex $\Delta^2 \subset \mathbb{R}^3$:

$$C = \{(\pi_R, \beta, \gamma) \in \Delta^2 : \beta, \gamma \in [a, b]; \pi_R + \beta + \gamma = 1\}$$

for some $[a, b] \subset [0, 1 - \pi_R]$. Keeping this in mind, the formula (3) becomes

$$\begin{aligned} U(w) &= \pi_R u(w_R) \\ &\quad + \alpha \min_{\beta \in [a, b]} [\beta u(w_G) + (1 - \pi_R - \beta)u(w_B)] \\ &\quad + (1 - \alpha) \max_{\gamma \in [a, b]} [\gamma u(w_G) + (1 - \pi_R - \gamma)u(w_B)] \end{aligned} \quad (4)$$

We interpret $a = b$ as absence of perceived ambiguity and $a = 0, b = 1 - \pi_R$ as maximal perceived ambiguity. 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.¹¹

To derive optimal choice behavior, it is convenient to work indirectly. Assume $a < b$ (so that the agent perceives ambiguity). Let $w = (w_R, w_G, w_B)$ be the optimal choice when prices are p . We begin by analyzing the implications of choices of relative wealth in the ambiguous states. There are three cases to consider:

- $w_G > w_B$: In this case, the minimum in the formula (4) for utility occurs when $\beta = a$ and the maximum occurs when $\gamma = b$, so the formula reduces to

$$\begin{aligned} U(w) &= \pi_R u(w_R) + \alpha a u(w_G) + \alpha(1 - \pi_R - a) u(w_B) \\ &\quad + (1 - \alpha)b u(w_G) + (1 - \alpha)(1 - \pi_R - b) u(w_B) \end{aligned}$$

Because we allow short sales, state wealth is not constrained to be positive, so the first order conditions for optimality are

$$\begin{aligned} \frac{\pi_R}{p_R} u'(w_R) &= \left[\frac{\alpha a + (1 - \alpha)b}{p_G} \right] u'(w_G) \\ &= \left[\frac{\alpha(1 - \pi_R - a) + (1 - \alpha)(1 - \pi_R - b)}{p_B} \right] u'(w_B) \end{aligned} \quad (5)$$

Rearranging the last equality yields

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

Because u' is strictly decreasing we obtain

$$\frac{p_G}{p_B} < \frac{(1 - \alpha)b + \alpha a}{\alpha(1 - \pi_R - a) + (1 - \alpha)(1 - \pi_R - b)} \quad (6)$$

¹¹This is a special implication of the fact that there are only two ambiguous states of nature.

- $w_G < w_B$: In this case, the minimum in the formula (4) for utility occurs when $\beta = b$ and the maximum occurs when $\gamma = a$, and a calculation similar to the above shows that

$$\frac{p_G}{p_B} > \frac{\alpha b + (1 - \alpha)a}{\alpha(1 - \pi_R - b) + (1 - \alpha)(1 - \pi_R - a)} \quad (7)$$

- $w_G = w_B$: In this case the minimum and the maximum in (4) are achieved for every choice of β, γ . Keeping the previous calculations in mind and doing a little algebra shows that the first-order conditions in this case are

$$\begin{aligned} \frac{p_G}{p_B} &\geq \frac{\alpha a + (1 - \alpha)b}{\alpha(1 - \pi_R - a) + (1 - \alpha)(1 - \pi_R - a)} \\ \frac{p_G}{p_B} &\leq \frac{\alpha b + (1 - \alpha)a}{\alpha(1 - \pi_R - b) + (1 - \alpha)(1 - \pi_R - a)} \end{aligned}$$

We can summarize the above discussion simply as:

$$w_G > w_B \iff \frac{p_G}{p_B} < \frac{(1 - \alpha)b + \alpha a}{\alpha(1 - \pi_R - a) + (1 - \alpha)(1 - \pi_R - b)} \quad (8)$$

$$w_G < w_B \iff \frac{p_G}{p_B} > \frac{\alpha b + (1 - \alpha)a}{\alpha(1 - \pi_R - b) + (1 - \alpha)(1 - \pi_R - a)} \quad (9)$$

$$w_G = w_B \iff \text{otherwise} \quad (10)$$

So far, we have made no assumption as to the agent's ambiguity aversion, but we now suppose that $\alpha \neq 1/2$, so that the agent is *not* maximizing expected utility. If $\alpha > 1/2$ — that is, the agent is ambiguity averse — then the right-hand side of (8) is strictly smaller than the right-hand side of (9). (Note that the numerators of the right-hand sides of inequalities (8) and (9) are convex combinations of a, b and that the denominators are convex combinations of $(1 - \pi_R - a), (1 - \pi_R - b)$.) Hence there is a *non-empty open set* of prices p for which an agent having these preferences will choose to hold an *unambiguous portfolio* — that is, a portfolio with $w_B = w_G$. Figure 1 illustrates the situation for $a = 0, b = 1 - \pi_R$. As indicated in the figure, with $\alpha > 1/2$ for $(\alpha, p_G/p_R)$ between the curves, the agent chooses $w_G = w_B$; for




Figure 1: Optimal choices as a function of α (with $a = 0$ and $b = 1 - \pi_R$)

$(\alpha, p_G/p_R)$ above the top curve, the agent chooses $w_G < w_B$; for $(\alpha, p_G/p_R)$ below the bottom curve, the agent chooses $w_G > w_B$. For α closer to 1, there is a larger range of prices for which the agent chooses an unambiguous portfolio; in particular, if $\alpha = 1$ (so the agent perceives maximal ambiguity and is maximally averse to ambiguity) the agent chooses an unambiguous portfolio for *all* possible prices.¹² Note that if $w_G = w_B$ is optimal — i.e., if condition (10) holds — then the following first-order condition will have to hold:

$$\frac{\pi_R u'(w_R)}{p_R} = \frac{(1 - \pi_R) u'(w_G)}{p_G + p_B} = \frac{(1 - \pi_R) u'(w_B)}{p_G + p_B} \quad (11)$$

¹²If $a > 0$ or $b < 1 - \pi_R$ then the qualitative features of Figure 1 remain, although the curves are closer together.

Finally, suppose that $\alpha < 1/2$, so the agent is ambiguity loving. Now the right-hand side of (8) is strictly *larger* than the right-hand side of (9), and there is an open set prices for which the agent has *two* solutions to her optimum problem, one with $w_G > w_B$, and one with $w_G < w_B$ (see Figure 1). For *no* price ratio p_G/p_B it is optimal for her to hold $w_G = w_B$. Thus, an ambiguity loving agent will *never* hold an unambiguous portfolio, and may for some price ratios be indifferent between a profile which pays more in state B than in state G and a profile which does the exact opposite. See Figure 1.

2.3 Equilibrium Implications

The implications derived above for individual choice have immediate implications for equilibrium choices and hence for equilibrium prices. Throughout, we assume $W_G \neq W_B$.

2.3.1 Homogeneous Ambiguity Attitudes

We first address the setting in which all agents have the same attitude toward ambiguity. It is convenient to discuss the various cases separately.

- **Case 1: all agents maximize expected utility with respect to a common prior $\pi = (\pi_R, \pi_G, \pi_B)$.**¹³ At equilibrium, all agents face the same prices and individual choices w^n sum to the social endowment $W = \sum e^n$, so it follows from (2) that

$$W_\sigma > W_\nu \implies \frac{p_\sigma}{\pi_\sigma} < \frac{p_\nu}{\pi_\nu} \quad \text{and} \quad w_\sigma^n > w_\nu^n \quad (12)$$

- **Case 2: $C^n = C$ and $\alpha^n = \alpha > 1/2$ for all n .** If $C = [0, 1 - \pi_R]$ and $\alpha = 1$, there is no equilibrium with positive prices, so we exclude this

¹³Recall that agents who are neutral to ambiguity — i.e., for whom $\alpha = 1/2$ — behave as if they maximize expected utility with respect to the uniform prior.

case.¹⁴ In other cases, 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 a common prior. If $W_G > W_B$, the imputed prior is

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

while if $W_G < W_B$, the imputed prior is

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

In either one of the above equilibria, we should observe all agents choosing portfolios which reflect the ranking of social wealth: if $W_G > W_B$ then $w_G^n > w_B^n$ for each n ; $W_G < W_B$ then $w_G^n < w_B^n$ for each n .

- **Case 3: $C^n = C$ and $\alpha^n = \alpha < 1/2$ for all n .** Here things are more complicated, because the optimization problem of an ambiguity loving agent may have *two* solutions, one with $w_G^n > w_B^n$, and one with $w_G^n < w_B^n$. If $W_G > W_B$, at an equilibrium some agents must settle on the $w_G^n > w_B^n$ choice but some might choose $w_G^n < w_B^n$. However, *no* agent will choose $w_G^n = w_B^n$, regardless of the price ratio p_G/p_B . Therefore, with common $\alpha < 1/2$ in equilibrium we should observe a group of agents whose holdings are ranked in the same direction as the aggregate wealth ratio, but we may also observe a group of agents whose holdings are ranked in the opposite direction from the aggregate wealth ratio. Since preferences are not convex, though, equilibrium is not guaranteed to exist in this case.

2.3.2 Heterogeneous Ambiguity Attitudes

We next turn to the setting in which attitudes toward ambiguity are heterogeneous across the population. As we shall see, equilibrium in this setting may be much more complicated. To illustrate, we suppose there are only

¹⁴There is an equilibrium in which some price is 0.

two types of agents: *Type I* agents maximize expected utility with respect to a common prior $\pi = (\pi_R, \pi_G, \pi_B)$; *Type II* agents perceive ambiguity (C is not a singleton) and are ambiguity averse ($\alpha > 1/2$). Write L for the set of Type I agents and M for the set of Type II agents.¹⁵ We suppose beliefs are consistent across the population, in the sense that $\pi \in C$. We are interested in the situation $W_G \neq W_B$; to be definite, assume $W_G > W_B$. (Of course, our conclusions change in the opposite way if $W_G < W_B$.)

Following our earlier discussion, we first derive implications for the distribution of equilibrium wealth in the ambiguous states. These implications are most easily expressed in terms of the distribution (across agents) of the share $w_G/(w_G + w_B)$ of wealth in state G relative to the total of wealth in the ambiguous states.

As (10) shows, the choices of Type II agents will depend on whether or not the equilibrium price ratio p_G/p_B falls in the interval

$$V = \left(\frac{(1 - \alpha)b + \alpha a}{\alpha(1 - \pi_R - a) + (1 - \alpha)(1 - \pi_R - b)}, \frac{\alpha b + (1 - \alpha)a}{\alpha(1 - \pi_R - b) + (1 - \alpha)(1 - \pi_R - a)} \right)$$

Notice that the interval V is increasing in C (perception of ambiguity) and in α (aversion to ambiguity). In the limit when $C = [0, 1 - \pi_R]$ and $\alpha = 1$, $V = (0, \infty)$, so the equilibrium price ratio p_G/p_B *necessarily* falls into V .

If p_G/p_B is *not* in V , then Type II agents will behave like Type I agents, and all agents will choose wealth holdings ranked in the same order as aggregate wealth. Hence all agents choose $w_G > w_B$ and $w_G/(w_G + w_B) > 1/2$. Moreover, there is no reason to expect qualitative differences between Type I and Type II agents. Note also that all agents will marginally adjust their holdings as the price ratio p_G/p_B changes.

If, on the other hand, the equilibrium p_G/p_B *is* in V then the situation

¹⁵ Similar results would follow if we assumed that all agents have α -maxmin preferences with common C , and Type I agents have strictly lower α than Type II agents, as long as equilibrium exists.

will be quite different.

- (i) Ambiguity averse 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 Type II agents. In particular, the population 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, Type I agents must hold, in aggregate, the remaining wealth; hence (with obvious notation):

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

The weighted average of the wealth shares $w_G^\ell/(w_G^\ell + w_B^\ell)$ of Type I agents must equal the average net wealth shares, we have:

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

Assuming that $W_a^{II} > 0$, recalling that we have assumed $W_G > W_B$ and doing a little algebra, we see that

$$W_G^I/(W_G^I + W_B^I) > W_G/(W_G + W_B)$$

That is, the distribution of wealth shares for Type I agents will be skewed to the right of the distribution of wealth shares that would be expected in the absence of ambiguity and ambiguity aversion.

- (iii) In view of (1), choices of Type I agents are sensitive to the entire vector p of state prices; in view of (10), (5) and (11), 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) Type II agents choose to hold equal wealth w_a^m in the ambiguous states G, B ; the magnitudes of w_a^m and wealth w_R^m in the risky state are determined by the budget constraints and first-order conditions (11). For these agents, state-dependent wealth need *not* be ranked opposite to the ranking of state/price probabilities. Hence the aggregate state-dependent wealth held by Type II agents also need *not* be ranked opposite to the ranking of state price/probabilities.
- (v) In the aggregate, Type I agents hold the difference between overall aggregate wealth and the aggregate wealth held by Type II agents. Because the ranking of state-dependent wealth held by Type II agents need not be ranked opposite to the ranking of state price/probabilities, it follows that the ranking of state-dependent wealth held by Type I agents need not be ranked opposite to the ranking of state price/probabilities either. However, as discussed before, the *individual* rankings of state-dependent wealth held by Type I agents should be the same as the *aggregate* ranking of state-dependent wealth held by Type I agents.

What rankings are possible? With respect to the ambiguous states, theory implies clear comparisons. Because $W_G > W_B$ and agents of Type II choose equal wealth in the ambiguous states G, B , $W_G^{II} = W_B^I$ and hence $W_G^I > W_B^I$. 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: $p_G/\pi_G < p_B/\pi_B$. However, no clear comparisons can be made with respect to the risky state R . As the reader can verify, 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 seem equally likely or plausible. For exam-

ple, simulations of environments in which each type consists of identical agents show that aggregate wealth rankings $W_G > W_R > W_B$ are less likely to lead to “anomalous” rankings of state price/probabilities than are aggregate wealth rankings $W_G > W_B > W_R$. In the latter case, the “anomalous” ranking $p_G/\pi_G < p_R/\pi_R < p_B/\pi_B$, which has a straightforward economic interpretation (the relative scarcity of Arrow security B and the high demand due to Type II agents shoots its probability-adjusted price higher than that of security R) is quite likely.

Finally, we note an implication for representative agent pricing. If all agents maximize expected utility with respect to a common prior (or more generally, have common ambiguity attitude and “beliefs”), then the ranking of state price/probabilities should be 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.¹⁶ However, 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 (state-independent) expected utility with respect to the common prior of Type I agents. Whether prices can be rationalized by the preferences of a representative agent with some *other* prior is an issue to which we shall return in Section 5.

¹⁶See 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. There were two endowment profiles (profile type i and profile type ii in the tables).¹⁷ 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.¹⁸ (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. Though some of the subjects had participated in previous economic experiments, no subject participated to more than one of our experiments.

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 \$2.00.¹⁹

¹⁷In 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.

¹⁸In 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.

¹⁹In the early sessions, we imposed this limit *ex post*, by barring subjects with more

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.²⁰

The state of the world was determined by a draw from an urn. At the beginning of each session, 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 —à la Ellsberg— 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. In particular, in sessions in which the composition of the urn was ambiguous, the ambiguity persisted throughout the session.

Sessions typically lasted 2.5 hours and began with two practice periods.²¹ Subject earnings ranged from \$0 to \$125, with an average of approximately \$50.

We classify experimental sessions according to the endowment distribution, the urn composition, and the ambiguity/risk environment. Following the discussion of discussed in Subsection 2.3, we used two endow-

than \$2.00 losses in a period 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.

²⁰Three 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.

²¹In some sessions subjects were paid in practice periods and in some sessions 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.

ment distributions, one chosen to make reversals in the ranking of state price/probabilities more likely (we refer to this treatment as PRR = possible/more likely rank reversals), and the other chosen to make rank reversals less likely (NRR = no/less likely rank reversals). For each endowment distribution we conducted experimental sessions with three different urn compositions (A, B, 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 risk and ambiguity and environments with pure risk.²² 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 endowment profile but *not* the endowments of others or aggregate endowments. In particular, subjects had no way of knowing the ranking of social wealth, and so could not distinguish between the NRR and PRR treatments.

The Appendix contains the web instructions that subjects read before starting an experimental session. Details of the last experimental session, classified as (NRR,B, Risk), can be viewed on the experimental web site: <http://clef.caltech.edu/exp/amb/start.htm>

²²It is worth repeating that no subject was confronted with *both* an ambiguous urn and a risky urn, since no subject participated to more than one experimental session.

Table 1: Security Endowments

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

Table 2: Initial Wealth

| Experiment Type | Profile 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 per Profile Type (i,ii)

| 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 state 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. The starkest 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 (PRR, 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 aggregate wealth ratios $W_G/(W_G + W_B)$ (which are approximately .98 in the NRR treatments and .63 in the PRR treatments; see Table 4). However, in the four Ambiguity treatments the modes have shifted to .50, apparently reflecting choices of ambiguity averse subjects, and the distributions have significantly bigger right tails, reflecting the compensating choices of ambiguity-tolerant subjects. (The few observations in the left tails — below 0.5 — are not strictly compatible with our simplified model, but they would be compatible with a small extension that allows for ambiguity loving agents because, as we discussed in Subsection 2.2, such agents *may* want to hold a portfolio with $w_G < w_B$.)




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.

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 our heterogeneous ambiguity attitude model, with the mode at .50 and heavy right tails.




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).

Table 6 further describes end-of-period holdings, confirming that 0.5 *is* the mode of relative holdings of ambiguous securities in each session (with ambiguity), and listing the proportion of subjects holding exactly the mode at the end of a period in each session. The last two rows present the second most frequent ratio after 0.5, and the proportion of subjects holding such ratio.²³ It is apparent that the differences in proportion are significant.

As the discussion of equilibrium in Section 2.3 shows, these holdings data are *not* compatible with any *homogeneous* ambiguity attitudes: if ambiguity attitudes were homogeneous across the population (in particular, if all agents maximized expected utility), we should *not* observe a significant number of agents holding a portfolio with $w_G/(w_G + w_B) = 1/2$.

²³In some (but not all) sessions, this second mode corresponded to the proportion induced by the initial endowments. The small numbers of subjects holding the second mode suggests the absence of endowment effects in our experiments.

Table 6: Sample Mode of Relative Holdings of Ambiguous Securities, Fraction of Subjects Holding It

| | NRR | | | PRR | | |
|-----------------------|------|------|------|------|------|------|
| | A | B | C | A | B | C |
| mode | 0.50 | 0.50 | 0.50 | 0.50 | 0.50 | 0.50 |
| % subj. | 0.16 | 0.13 | 0.18 | 0.24 | 0.09 | 0.18 |
| 2nd mode ^a | 0.59 | 0.83 | 1.00 | 0.30 | 0.64 | 0.67 |
| % subj., 2nd mode | 0.01 | 0.01 | 0.02 | 0.02 | 0.03 | 0.01 |

^aMode of distribution after eliminating the above mode.

To explain further, note first that for an agent who maximizes expected utility, a portfolio with the proportion $w_G/(w_G + w_B) = 1/2$ is only optimal if

$$\frac{p_G}{\pi_G} = \frac{p_B}{\pi_B} \quad (13)$$

If all agents maximize expected utility and hold common priors, this condition *cannot* hold for all agents because the ranking of social endowments is $W_G > W_B$. If all agents maximize expected utility, but perhaps with different priors, then (13) may hold for agents with particular priors — but it requires a knife-edge condition on priors; hence we would not expect to observe it for a significant number of agents. In either case: if all agents maximized expected utility (even with different priors) we would not expect to observe histograms with a spike on the proportion 1/2.

Similarly, the discussion of Section 2.3 shows that a spike on the proportion 1/2 is incompatible with homogeneous aversion to ambiguity ($\alpha > 1/2$). It is also incompatible with homogeneous love of ambiguity ($\alpha < 1/2$), because ambiguity loving agents should choose a proportion $w_G/(w_G + w_B)$ that is either strictly above 1/2 or strictly below 1/2, but never equal to 1/2.

Heterogeneity in ambiguity aversion has additional implications. As already pointed out when discussing Figure 1, no matter how prices change, an agent with extreme ambiguity aversion ($\alpha = 1$) will choose not to be

exposed to ambiguity, and hence, hold $w_G/(w_G + w_B) = 0.5$. As ambiguity tolerance increases, agents may continue not to be exposed to ambiguity, but the range of prices for which this obtains shrinks. Now, across periods in our experiments, prices did change. With changing prices, we therefore expect variation of $w_G/(w_G + w_B)$ to be higher for more ambiguity-tolerant agents. Identifying ambiguity tolerance by average exposure to ambiguity (average deviation of $w_G/(w_G + w_B)$ from 0.5 across periods), we therefore predict the variability of subjects' exposure to ambiguity (standard deviation of $w_G/(w_G + w_B)$ across periods) to increase with average exposure to ambiguity. Figure 5 confirms this prediction, lending further support to heterogeneous ambiguity attitudes model.




Figure 5: Plot of variability (standard deviation across periods) of wealth allocated to state G as a proportion of final wealth in states G and B (the ambiguous states), against average deviation (across periods) of this proportion from 0.5; each observation corresponds to one subject in one (ambiguity) experiment.

The expected utility model of Savage (1954) would be consistent with the findings in Figure 5 only if subjects who happened to have beliefs in one period that make them hold close to equal amounts of the ambiguous state securities also changed their beliefs less, which would require more obstinate priors. That is, explanation of Figure 5 in terms of the expected utility model requires one to argue that precision of priors is correlated with exposure to the ambiguous states (which, incidentally, the expected utility agents do not perceive as ambiguous).

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. 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 possible, but they do not yield uniformly better results; see Section 5 for discussion.)

We emphasize 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 experimental sessions, 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 between the predictions of a benchmark model with homogeneous (or nil; i.e., expected utility) ambiguity aversion and a our model with heterogeneous ambiguity attitudes.

First, consider the NRR treatment. By construction, $W_G > W_R > W_B$, so in the Risk treatment the benchmark model predicts $p_B/\pi_B > p_R/\pi_R > p_G/\pi_G$. Moreover, as the discussion in Subsection 2.3 suggests, our model predicts that the same ordering should be most likely in the Ambiguity treatment as well. As Figure 6 shows, this is what we see in the data. (The top panels of Figure 6 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 .

Next, consider the PRR treatment, where by construction, $W_G > W_B > W_R$. As we have discussed in Section 2, in all the Risk treatments the benchmark model predicts $p_R/\pi_R > p_B/\pi_B > p_G/\pi_G$. However, in the Ambiguity treatments our model suggests that we may see rank reversals, likely leading to the ordering $p_B/\pi_B > p_R/\pi_R > p_G/\pi_G$. As the left panels of Figure 7 show, for the A session 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 B sessions the data speak less clearly, and in the Risk version, we see anomalous rankings: the ECDF for




Figure 6: 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 .

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. In later periods, subjects seemed to believe (perhaps because of a belief in the “law of small numbers”) that 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.²⁴ In the corresponding Ambiguity treatment, the ECDF for B is shifted upward and very close to the ECDF

²⁴It is not clear that this kind of problem can be avoided. Of course one could exercise some control over the sequence of draws — but then the draws would no longer be random.




Figure 7: 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 .

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. The data for the final two sessions are shown in Figure 8. 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 or everybody was equally ambiguity averse. In the right panel, for which the configuration is (PRR, C, Ambiguity), the predicted rankings under homogeneous ambiguity aversion would be $p_R/\pi_R > p_B/\pi_B > p_G/\pi_G$; in the actual data the rankings appear anomalous.




Figure 8: 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 .

As we mentioned earlier, we emphasize the prices for *all* the transactions because it is difficult — if not impossible — to assess when prices have “settled down” during an experimental period. However, it is interesting to consider what happens to the trade prices as the experimental session progresses; i.e., as periods go by. Table 7 reports the, per period, per experiment (excluding the initial two practice periods) averages of two ratios:

- for the NRR treatments: the average of $(p_B/\pi_B)/(p_R/\pi_R)$; averages marked ** have significance above 99%
- for the PRR treatments: the average of $(p_R/\pi_R)/(p_B/\pi_B)$; averages marked oo have significance above 99%

As recalled above, under homogeneous ambiguity aversion all the ratios in the table should be strictly greater than 1. The table shows that this *did* obtain in later periods in the NRR experiments, but that the opposite — what is likely to happen according to our model— was observed in the PRR experiments. In those experiments, as periods progressed the ratio tended to

Table 7: Average of $(p_B/\pi_B)/(p_R/\pi_R)$ (NRR) and $(p_R/\pi_R)/(p_B/\pi_B)$ (PRR), per Period, per Experiment

| Treatment | Experiment | Period Number | | | | | |
|-----------|------------|---------------|--------|--------|--------------------|--------------------|--------------------|
| | | 1 | 2 | 3 | 4 | 5 | 6 |
| NRR | A | 0.96 | 1.27** | 1.61** | 2.02** | 2.47** | 1.88** |
| | B | 1.26** | 1.17** | 1.52** | 1.34** | 1.48** | 1.71** |
| | C | 1.19** | 1.49** | 1.50** | 2.25** | 2.61** | 2.97** |
| PRR | A | 1.14** | 1.10** | 1.07 | 0.94 ^{oo} | 0.65 ^{oo} | 0.76 ^{oo} |
| | B | 1.00 | 0.98 | 1.05 | 0.97 | 0.88 ^{oo} | 0.50 ^{oo} |
| | C | 1.70** | 1.36** | 1.37** | 1.27** | 0.97 | 0.50 ^{oo} |

“settle” in favor of the ranking $p_B/\pi_B > p_R/\pi_R$. We therefore see that the “anomalous” price rankings appear clearly in the late periods of each of the PRR experiments, even when they are not highlighted by the ECDF plots.

Summing up, the pricing effects due to the introduction of ambiguity are consistent with the suggestions of the theoretical analysis of Section 2: rank changes in state price/probability ratios are observed, but only in the PRR treatment —in which the security in shortest supply pays off in a risky, rather than ambiguous, state of the world.

As we argued earlier, the heterogeneous ambiguity model we suggest has the feature that a fraction of agents — those whose are ambiguity averse — is *infra-marginal* for the price ratio p_B/p_G : they hold $w_B^m = w_G^m$ for an interval of price ratios (or perhaps for all price ratios, if they are sufficiently ambiguity averse). The remaining agents, who are less ambiguity averse or ambiguity neutral, are the only ones who are marginal for the price ratio p_B/p_G . However, it is *not* true that the presence of the most ambiguity averse agents does not affect the price ratio. It does not do so *directly*, of course, but it does so *indirectly*: as the proportion of ambiguity averse agents

grows, the imbalance in the relative supplies of ambiguous securities has to be borne by fewer and fewer agents, who demand more attractive prices.




Figure 9: The effects of ambiguity averse (i.e., who hold approximately $w_G = w_B$) agents on p_B/p_G , with linear fits.

Indeed, that is precisely what we observe in our experimental data. Figure 9 shows that as the proportion of agents with close to $w_G = w_B$ end-of-period holdings in a given experimental period increases,²⁵ the average ratio p_B/p_G observed in the period increases as well, both in the NRR (left) and in the PRR (right) treatments. The figure also reports linear fits, which are jointly significant at the 10% confidence level.

4.3 Risk Aversion and Ambiguity Aversion

In theory, there seems to be no reason why risk aversion — in our framework, concavity of the felicity function u — and ambiguity aversion — in our frame-

²⁵Precisely, these are: for the NRR experiments the subjects whose end-of-period ratio $w_G/(w_G + w_B)$ is within 0.1 of 0.5 (compare to market portfolio proportion, which is 0.98); for the PRR experiments the subjects whose end-of-period ratio $w_G/(w_G + w_B)$ is within 0.025 of 0.5 (compare to market portfolio proportion, which is 0.63).

work, the coefficient α — should be at all correlated, but our experimental data suggests that they may in fact be positively correlated.

To see this, 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 10 displays the average of such ranges for all periods in all the sessions what involved ambiguity. We observe 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). Agents who are close to ambiguity neutrality are also almost risk neutral.²⁶

A significant positive correlation between ambiguity aversion and risk aversion implies a particular kind of market segmentation — particular kinds of assets are held disproportionately by less risk-averse individuals — and therefore has substantial implications for asset pricing. It suggests, 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 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, and the resulting market segmentation, might also be relevant for regulation (Easley & O’Hara, 2005).

²⁶Our findings are consistent with at least one study in neuroscience (Hsu *et al*, 2005).




Figure 10: 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 average range of wealth allocated across all states; all periods in all experiments with ambiguous states.

5 Discussion

There are two issues that deserve more discussion. The first is connected to the previously mentioned question of which prior should be used in our Ellsberg setting to deflate the observed state prices and the implications of that choice for representative agent pricing. The second is whether a different model of ambiguity aversion might be better suited for explaining our data.

As we remarked earlier, our discussion of the pricing results of our experiments used a uniform initial probability over ambiguous states (with subsequent Bayesian updating) in the calculation of state price/probabilities. As we have already observed, this choice seems to fit well with the data for the

NRR treatments. In particular, if all agents maximize expected utility and Bayesian update from such a prior — or if agents entertain different priors, but there is a “representative” agent with such a prior — we should observe state price/probabilities similar to those in the data. However, in the PRR treatments, this choice would *not* fit the data well, especially not in the later periods of experimental sessions.

It is true, however, that for each *particular* experimental session, it would be possible to find an *ad hoc* prior for which the ranking of the ratios is not “perverse.” Put differently, for each *particular* experimental session, it would be possible to find an *ad hoc* prior for which pricing is consistent with the existence of a representative agent. Indeed, it seems that data from a *single* experimental session would *always* be consistent with existence of a representative agent provided we choose the priors of that agent carefully enough. What does *not* seem to be possible is to find a *single* prior that delivers “non-perverse” rankings for *all* the sessions. Based on our simulations, the uniform prior is the one which delivers the fewest “perverse” rankings, and therefore the most favorable to a strictly “common prior” Bayesian model.

One could conjecture that in our Ellsberg setting it would be natural to expect agents to entertain different priors, and that equilibrium forces might lead the market to behave “as if” there were a representative agent with state-independent utility but with prior beliefs that differed across sessions in a way that depended on how the “real” beliefs were distributed among the agents.²⁷ This would explain the need to use different deflators in different experiments. However, to be persuasive, such a conjecture would need to answer (at least) two questions: 1) Why should the distribution of prior beliefs display patterns which depend on experimental conditions (the aggregate endowments) of which the agents are *not informed*? In particular, why should the “representative” prior be the uniform prior in the NRR sessions but different in the different PRR sessions? 2) More importantly, as discussed earlier, a Bayesian agent will choose to have equal wealth in the ambiguous

²⁷For an example of a representative agent theorem under heterogeneous beliefs, see Jouini & Napp (2007).

states ($w_G = w_B$) only in the knife-edge condition in which 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. This is especially true for the NRR sessions, in which the social endowments W_G, W_B and prices p_G, p_B are quite far apart. 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. Thus, we think that heterogeneity in ambiguity attitude is the driving force behind our experimental observations, rather than heterogeneity in beliefs.

The significant proportion of agents holding equal wealth in the ambiguous states is also the reason why we do not think it appropriate to model ambiguity averse agents via the so-called “smooth ambiguity” model of Klibanoff, Marinacci & Mukerji (2005). In that model, risk aversion and ambiguity aversion have qualitatively similar implications: an ambiguity averse agent will choose to hold equal wealths $w_G = w_B$ in the ambiguous states only for a single price ratio p_G/p_B , and will readjust his holdings in the ambiguous states as soon as that ratio changes. Again, the end-of-period holdings we observe in our data suggest that a significant number of agents choose to hold equal wealths in the ambiguous states for a range of price ratios; this behavior seems more consistent with the α -maxmin model we use.²⁸

²⁸These results have recently been confirmed by Ahn, Choi, Gale & Kariv (2009) in an experiment on individual portfolio choice.

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 are at odds with two apparently wide-spread and often-asserted beliefs. The first is that prices reflect an average of the beliefs of *all* agents.²⁹ 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.

²⁹See Hirshleifer (2001) for instance.

References

- Ahn, D, S. Choice, D. Gale and S. Kariv (2009): “Estimating Ambiguity Aversion in a Portfolio Choice Experiment,” UC Berkeley working paper.
- 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 (2006): “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 (2009): “Ambiguity And Nonparticipation: The Role of Regulation,” *Review of Financial Studies*, forthcoming.
- 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.
- Izhakian, Y. and S. Benninga. (2008): "The Uncertainty Premium in an Ambiguous Economy." Tel-Aviv University working paper.
- Jouini, E. and C. Napp (2007): "Consensus Consumer and Intertemporal Asset Pricing with Heterogeneous Beliefs." *Review of Economic Studies* 74, 1149-1174.

- Klibanoff, P., Marinacci, M. and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity." *Econometrica* 73, 1849-1892.
- 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* 59, 969-997.
- Knight, Frank (1939): *Risk, Uncertainty and Profit*, London: London School of Economics.
- Maenhout, P. (2004): "Robust Portfolio Rules and Asset Pricing." *Review of Financial Studies* 17, 951-983.
- Savage, L.J. (1954): *The Foundations of Statistics*, J. Wiley and Sons, New York.
- Skiadas, C. (2008): "Dynamic Portfolio Choice and Risk Aversion," in: *Handbooks in OR&MS Volume 15*, J.R.Birge and V. Linetsky, eds., Elsevier, Chapter 19.
- Trojani, F., M. Leippold and P. Vanini (2007): "Learning and Asset Prices under Ambiguous Information," *Review of Financial Studies* 21, 2565-2597.
- Uppal, R. and T. Wang (2003): "Model Misspecification and Under Diversification." *Journal of Finance*, 58, 2465-2486

Instructions

Overview:

- I. [The Experiment](#)
- II. [The Markets Interface, jMarkets](#)

I. THE EXPERIMENT

1. Situation

The experiment consists of a sequence of trading sessions, referred to as *periods*. At the beginning of a period, you will be given a fresh supply of *securities* and *cash*. Markets open and you are free to trade your securities. You buy securities with cash and you get cash if you sell securities.

At the end of the period, after markets close, the securities expire, after paying *dividends* that will be specified below. These dividends, together with your cash balance, constitute your *period earnings*.

Period earnings are *cumulative* across periods. At the end of the experiment, the cumulative earnings are yours to keep, in addition to a standard sign-up reward.

During the experiment, accounting is done in real dollars.

2. The Securities

You will be given two types of securities, *stocks* and *bonds*. Bonds pay a fixed dividend at the end of a period, namely, \$0.50. Stocks pay a random dividend. There are three types of stocks, referred to as X, Y and Z. Their payoff depends on the drawing from an urn, as explained later. The payoff is either \$0.50 or nothing. When X stock pays \$0.50, Y and Z stock pay nothing; when Y stock pays \$0.50, X and Z stock pays nothing; when Z stock pays \$0.50, X and Y stock pay nothing.

You won't be able to *buy* stock or bonds unless you have the cash. You will be able to *sell* stock and bonds (and get cash) even if you do not own any. This is called *short selling*. If you sell, say, one X stock, then you get to keep the sales price, but \$0.50 will be subtracted from your period earnings after the market closes and if the payoff on X stock is \$0.50; you don't pay anything if the payoff on X is zero. If at the end of a period you are holding, say, -1 bonds, \$0.50 will be subtracted from your period earnings.

The trading system checks your orders against *bankruptcy*: you will not be able to submit orders which, if executed, are likely to generate losses of more than \$2 at the end of the period.

3. How Payoffs Are Determined

There are three possible states, X, Y and Z. Stock X pays when state X occurs; stock Y pays when state Y occurs; stock Z pays when state Z occurs.

| | Stock X | Stock Y | Stock Z |
|---------------|---------|---------|---------|
| If state is X | \$0.50 | \$0 | \$0 |
| If state is Y | \$0 | \$0.50 | \$0 |
| If state is Z | \$0 | \$0 | \$0.50 |

Here is how states are drawn. Imagine an urn with, say, 20 balls, 7 of which are marked X (X balls), 7 are marked Y (Y balls), and 6 are marked Z (Z balls). To determine the state in the first period, we draw one ball from this urn. Imagine that we draw an X ball. This will determine the payoff on the stocks: for each unit of stock X you're holding at the end of the period, you will receive \$0.50. You will not receive anything for your holdings of stocks Y and Z. We then throw away this ball. That is, the X ball is not placed back in the urn. As a consequence, we draw the state for period 2 from an urn with the following composition: 6 X balls, 7 Y balls, and 6 Z balls. If we draw a Z ball in the second period, then the state for the third period will be drawn from an urn with the

following composition: 6 X balls, 7 Y balls, and 5 Z balls. Etc.

We draw the state at the beginning of each period. This means that the state will not depend on what you do during the period. Nobody will be told what the state is until the end of the period.

The initial composition of the urn is announced in the *News* page.

II. THE MARKETS INTERFACE, jMARKETS

Once you click on the *Participate* link to the left, you will be asked to log into the markets, and you will be connected to the *jMarkets* server. After everybody has logged in and the experiment is launched, a markets interface like the one below will appear.

3 of 6

21.01.2008 16:06

Experiment Instructions

http://clef.caltech.edu/exp/amb/instructions.html

JMarkets Client Ten Markets

Ten Markets (Subject ID: 1)

Time Remaining: 169
Total Cash Holdings: \$1.00

Active Markets

| Stock X | Stock Y | Stock Z | Bond |
|--|--------------|--------------|--------------|
| Holdings: 0 | Holdings: 12 | Holdings: 20 | Holdings: -5 |
| Center | Center | Center | Center |
| \$0.17 | \$0.50 | \$0.50 | \$0.50 |
| Buy (0) Sell (2) \$0.16 Your offers: 1 Cancel | \$0.49 | \$0.49 | \$0.49 |
| \$0.15 | \$0.48 | \$0.48 | \$0.48 |
| \$0.14 | \$0.47 | \$0.47 | \$0.47 |

History

Price Chart

Current Orders

| Security | Price | Offer | Units |
|----------|-------|------------|-------|
| Bond | 0.48 | Buy Order | 1 |
| Bond | 0.5 | Sell Order | 1 |
| Stock X | 0.14 | Buy Order | 1 |
| Stock X | 0.15 | Buy Order | 1 |
| Stock X | 0.16 | Sell Order | 1 |

Earnings History

| Period | Security | Final Holdings | Cumulative E |
|--------|----------|-----------------|--------------|
| 0 | Stock X | 0 (+3.0, -3.0) | |
| 0 | Stock Y | 12 (+0.0, -0.0) | |
| 0 | Stock Z | 20 (+0.0, -0.0) | |
| 0 | Bond | -5 (+0.0, -0.0) | |

4 of 6

21.01.2008 16:06

1. Active Markets

The *Active Markets*

panel is renewed each period. In it, you'll see several scroll-down columns. Each column corresponds to a market in one of the securities. The security name is indicated on top. At the bottom, you can see whether the market is open, and if so, how long it will remain open. The time left in a period is indicated on the right hand side above the Active Markets panel.

At the top of a column, you can also find your current holdings of the corresponding security. Your current cash holdings are given on the right hand side above the Active Markets panel.

Each column consists of a number of price levels at which you and others enter offers to trade. Current offers to sell are indicated in red; offers to buy are indicated in blue. When pressing the *Center* button on top of a column, you will be positioned halfway between the best offer to buy (i.e., the highest price at which somebody offers to buy) and the best offer to sell (i.e., the lowest price that anybody offers to sell at).

When you move your *cursor*

to a particular price level box, you get specifics about the available offers. On top, at the left hand side, you'll see the number of units requested for purchase. Each time you click on it, you send an order to buy one unit yourself. On top, at the right hand side, the number of units offered for sale is given. You send an order to sell one unit each time you yourself click on it. At the bottom, you'll see how many units you offered. (Your offers are also listed under *Current Orders* to the right of the Active Markets panel.) Each time you hit *cancel*, you reduce your offer by one unit.

If you click on the price level, a small window appears that allows you to offer multiple units to buy or to sell, or to cancel offers for *multiple* units at once.

2. History

The *History* panel shows a chart of past transaction prices for each of the securities. Like the *Active Markets* panel, it refreshes every period.

3. Current Orders

The *Current Orders*

panel lists your offers. If you click on one of them, the corresponding price level box in the Active Markets panel is highlighted so that you can easily modify the offer.

4. Earnings History

The *Earnings History* table shows, for each period, your final holdings for each of the securities (and cash), as well as the resulting period earnings.

5. How Trade Takes Place

Whenever you enter an offer to sell at a price below or equal to that of the best available buy order, a sale takes place. You receive the price of the buy order in cash. Whenever you enter an offer to buy at a price above or equal to that of the best available sell order, a purchase takes place. You will be charged the price of the sell order.

The system imposes strict price-time priority: buy orders at high prices will be executed first; if there are several buy orders at the same price level, the oldest orders will be executed first. Analogously, sell orders at low prices will be executed first, and if there are several sell orders at a given price level, the oldest ones will be executed first.

6. Restrictions On Offers

Before you send in an offer, *jMarkets* will check two things: the *cash constraint*, and the *bankruptcy constraint*.

The cash constraint concerns whether you have enough cash to buy securities.

If you send in an offer to buy, you need to have enough cash. To allow you to trade fast, *jMarkets* has an automatic cancelation feature. When you submit a buy order that violates the cash constraint, the system will automatically attempt to cancel buy orders you may have at lower prices, until the cash constraint is satisfied and your new order can be placed.

The bankruptcy constraint concerns your ability to deliver on promises that you implicitly make by trading securities. We may not allow you to trade to holdings that generate losses in some state(s). A message appears if that is the case and your order will not go through.

Good Luck!