

Prices and Allocations in Asset Markets with Heterogeneous Attitudes Towards Ambiguity

Peter Bossaerts, Paolo Ghirardato, Serena Guarnaschelli and William Zame[‡]

This version: February 2007

[‡]Bossaerts: Caltech, Université de Lausanne (Swiss Finance Institute) and CEPR; Guarnaschelli: McKinsey; Ghirardato: Università di Torino & Collegio Carlo Alberto; Zame: UCLA.

Prices and Allocations in Asset Markets with Heterogeneous Attitudes Towards Ambiguity

Abstract

This paper studies the impact of ambiguity aversion on equilibrium asset prices and portfolio holdings in competitive financial markets. It argues that attitude toward ambiguity is heterogeneous in the population, just as attitude toward risk is heterogeneous in the population, but that heterogeneity in attitude toward ambiguity has different implications than heterogeneity in attitude toward risk. Specifically, agents who are sufficiently ambiguity averse find *open* sets of prices for which they refuse to hold an ambiguous portfolio. This leads to a wider range of state price densities and to potential reversals of ranking of state price-probability ratios relative to aggregate wealth. In addition, the distribution of holdings will have a new mode, with highly ambiguity averse agents holding securities with ambiguous payoffs in equal proportions. Under pure risk, the distribution of holdings has a single mode equal to the market portfolio weight. Experiments confirm the theoretical predictions. While price patterns often look little different from those under pure risk, portfolio choices display strong effects from the presence of ambiguity. The experiments also suggest a positive correlation between risk aversion and ambiguity aversion, which may explain the “value effect” in field data.

Prices and Allocations in Asset Markets with Heterogeneous Attitudes Towards Ambiguity[†]

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 individual agents distinguish between *risk* (known probabilities) and *ambiguity* (unknown probabilities), and that individual agents 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 combine theory and experiment to show that ambiguity and ambiguity aversion can have a significant impact on equilibrium asset prices and portfolio holdings in competitive financial markets. We find that attitude toward ambiguity is heterogeneous in the population, just as attitude toward risk is heterogeneous in the population, but that heterogeneity in attitude toward ambiguity has different implications than heterogeneity in attitude toward risk. Agents who are risk averse will refuse to hold a risky portfolio only if price ratios are exactly equal to ratios of expected payoffs, which is a knife-edge condition, but agents who are ambiguity averse will refuse to hold an ambiguous portfolio for an *open* set of prices. This ambiguity aversion may lead to reversals of ranking of state price-probability ratios relative to aggregate wealth, and hence, to asset pricing that is *inconsistent* with the existence of a representative agent who maximizes state-independent utility. The heterogeneity in ambiguity aversion generates bi-modal distributions of portfolio shares, with highly ambiguity averse agents holding ambiguous (state) securities in equal proportion.

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 red, blue and green balls. One third of the balls in the urn are known to be red; the remaining balls are blue or green, in unknown proportion. Ellsberg asked subjects first, whether

[†]We are grateful for comments to seminar audiences at CIRANO, U.C. Irvine, Kobe University, Collegio Carlo Alberto, Sankt Gallen University, the Securities and Exchange Commission, the 2005 Atlanta Fed Conference on Experimental Finance, the 2006 Skinance Conference in 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), National Science Foundation grants SES-0079374 (Bossaerts) and SES-0079374 (Zame), the UCLA Academic Senate Committee on Research (Zame), and the Swiss Finance Foundation (Bossaerts). Views expressed here are those of the authors and do not necessary reflect the views of any funding agency.

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

²Exceptions include Epstein and Wang (1994) and Cagetti e.a. (2002).

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 (and experimentalists since) found that many subjects prefer “red” in each of the the first choices and “blue or green” in the second. Such behavior is “paradoxical” — in the sense that it is inconsistent with maximizing expected utility with respect to any subjective prior, and hence violates the Savage (1954) axioms (specifically, the “sure thing principle”).

In our environment, uncertainty is again identified with a draw of a ball from an urn that contains red, blue and green balls. Initially, the total number of balls is known, and a given number of balls are known to be red; the remaining balls are blue or green, in unknown proportions. Draws from the urn are made without replacement. (Thus, agents must update after each draw.) Before each draw from the urn, subjects are endowed with and can trade Arrow securities, each of which pays the same fixed amount according to the draw of a particular color. Note that the red security (i.e., the security that pays when a red ball is drawn) is risky (that is, the distribution of its payoffs is known) and that the blue and green securities are ambiguous (that is, the distribution of their payoffs is unknown). We study the effect of ambiguity and ambiguity aversion and the relative supply of risky and ambiguous securities on equilibrium prices and holdings.^{3,4}

To model preferences that display ambiguity aversion, we use the “ α -maxmin” model of Ghirardato e.a. (2004), which is a generalization of the “maxmin” model of Gilboa and Schmeidler (1989). This specification provides a natural way to broaden the spectrum of agents’ behavioral traits; it does so without a radical departure from the familiar expected utility model and with little loss in terms of tractability. As we show, the crucial property of α -maxmin preferences is that there is an open interval containing 1 such that if the ratio of prices of the ambiguous securities lies in this interval then an agent with these preferences will always choose to hold an unambiguous portfolio — that is, a portfolio containing equal quantities of the ambiguous securities. For the extreme case $\alpha = 1$ (greatest ambiguity aversion), an agent with those preferences 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 (non-zero) quantities of two securities only if the ratio of prices is equal to the ratio of subjective state probabilities.

Our experimental data clearly show that some agents are quite ambiguity averse. However, the same data show that the population is heterogeneous: some agents appear less ambiguity averse than others, and some agents appear to maximize expected utility.

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 blue or green security should be higher than the price of the red or blue security. Of course, such prices could not obtain at a market equilibrium because they admit an arbitrage opportunity.

⁴Epstein and 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.

Heterogeneity in ambiguity attitudes has important effects. In a market populated by agents who maximize expected utility with respect to identical priors (but perhaps different risk attitudes), equilibrium prices can always be rationalized by a representative agent who maximizes expected utility (with respect to the same prior). In particular, state price-probability ratios (ratios of state prices to probabilities) are necessarily ordered in the opposite direction as supplies of state securities. That is, the state price-probability ratio will be highest in the state for which social endowment is least. However, heterogeneity in ambiguity attitudes opens the following possibility: If the ambiguous securities are in unequal total supply and the most ambiguity averse agents choose unambiguous portfolios, the imbalance of ambiguous securities must be held by the least ambiguity averse agents. This effect may distort the state price-probability ratios.⁵ If the distortion is small, state price-probability ratios will still be ordered in the direction opposite to supply; if the distortion is greater, state price-probability ratios may not be ordered in the direction opposite to supply. In our experiments, we manipulate supplies so that, in one treatment, anomalous orderings are unlikely to occur, and in another treatment, anomalous orderings are possible. Correspondingly, we record no state price-probability ranking violations in the first treatment, but observe violations in the majority of experiments in the second treatment.

When state price densities are not ordered in the direction opposite to supply, equilibrium prices *cannot* be rationalized by a representative agent who maximizes expected utility or even by a representative agent who is ambiguity-averse. In such a case, equilibrium prices can only be rationalized by a representative agent with state-dependent utility. This finding would seem to have important implications for finance, where the representative agent methodology is pervasive.

Heterogeneity in ambiguity attitudes has another effect. Because the more ambiguity averse agents choose to hold an unambiguous portfolio, other agents must hold ambiguous securities in more extreme proportions in a setting with ambiguity than they would in a setting with pure risk. Thus, the cross-sectional variation of portfolio holdings should be more extreme when there is ambiguity. In practice, under pure risk, the distribution of holdings has a single mode equal to the market portfolio weight. Under ambiguity, however, we predict that a second mode should appear, because highly ambiguity averse agents hold ambiguous state securities in exactly equal proportions. Others should compensate, forcing the first mode to move away from the market portfolio weight.

We confirm these holdings predictions. To facilitate the interpretation of subjects' choices under ambiguity, we replicated many experiments eliminating ambiguity. In these replications, subjects were told the exact probabilities of all of the states. Distributions of portfolio weights indeed exhibit unique modes equal to market portfolio weights. These choices obtain despite strong bi-modality of initial allocations. In contrast, under ambiguity, distributions are more spread out, and there is always a new mode reflecting equal weighting of ambiguous state securities. At times,

⁵In the presence of ambiguity, we define state price-probability ratios using a uniform prior over ambiguous states. We discuss alternatives later in the paper.

the first mode moves away from the market portfolio weight.

In principle, there need be no correlation between ambiguity aversion and risk aversion. However, the 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 would be held — and priced — only by investors who are less ambiguity averse and hence 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.⁶

Like us, Easley and O’Hara (2005) 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. Unlike our paper, however, there is no empirical analysis to demonstrate that such an effect can be observed in practice. Instead, the paper focuses on the theoretical implications for regulation. In deriving such implications, the paper assumes that risk aversion and ambiguity aversion are uncorrelated. Some of the implications may be invalidated if, as we find, attitudes towards risk and towards ambiguity are correlated.

The present paper is related to a the recent literature in several ways. Most obviously, the studies here are an outgrowth of Bossaerts e.a. (2005), which studies environments with pure risk. Bossaerts, Plott & Zame document that there is substantial heterogeneity in preferences (demand functions) but that much of this heterogeneity washes out in the aggregate, so that representative agent pricing obtains. In the environment addressed here, with both risk and ambiguity, heterogeneity does not wash out in the aggregate and representative agent pricing does not obtain.

This paper adds to an emerging literature that uses experiments to study the effect of “irrational” preferences (eg., preferences other than expected utility) on prices and choices in competitive markets through experiments. Gneezy e.a. (2003) analyzes the impact of myopic loss aversion on pricing, but assumes homogeneous preferences. Kluger and Wyatt (2004) studies the impact of particular cognitive biases in updating on 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 and Polkovnichenko (2005) studies 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; in particular, equilibrium prices always have a representative agent rationalization. See also Fehr and Tyran (2005) for an overview.

⁶We thank Nick Barberis for pointing this out.

There is a recent literature which seeks to explain the equity premium puzzle (high average returns on equity and low average riskfree rate) by appealing to ambiguity (or, equivalently, Knightian or model uncertainty) on the basis of a model with an ambiguity-averse representative agent [Epstein and Wang (1994), Uppal and Wang (2003), Cagetti e.a. (2002), Maenhout (2000), Skiadas (2005), Trojani, Leippold and Vanini (2005)]. Our results suggest a problem with this literature, since heterogeneity in ambiguity attitudes leads (even in relatively small groups) to pricing that are not rationalizable with a representative agent.

Following this Introduction, Section 2 presents a theoretical analysis, generating asset pricing and portfolio holding predictions. 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 Theoretical Predictions

Absent ambiguity aversion (i.e., in the traditional context of expected utility preferences), the effects of heterogeneity in risk aversion on equilibrium prices and holdings is straightforward to derive, because all agents remain marginal; more risk averse agents merely reduce their exposure to risk. When ambiguity is allowed into the picture and the level of aversion to ambiguity differs across agents, the effects are less obvious. This is because some agents may not be marginal any more, effectively eliminating their influence on the relative pricing of ambiguous securities.

[I deleted the different cases and go straight to different levels of ambiguity - that's the goal]

Let there be three states $s = r, g, b$. Three state securities, also labeled $j = r, g, b$, are available, in addition to riskfree notes and cash (which are perfect substitutes, and therefore can be ignored). State security j pays 1 in state s where $s = j$ and zero otherwise. State $s = r$ is *risky* in the sense that it has less than unit probability of occurring, but this probability is known and equal to π . States $s = g$ and $s = b$ are *ambiguous*, in the sense that the conditional probability of $s = g$ given either $s = g$ or $s = b$ is unknown. The probability of “either $s = g$ or $s = b$ ” is, however, known and equal to $1 - \pi$.

2.1 Optimal Investments

We first discuss optimal investment for expected utility preferences, and then move to preferences with sensitivity to ambiguity. With juxtapose these two types of preferences in order to highlight the differences.

Suppose that subject $n \in \{1, \dots, N\}$ has (state-independent) expected utility preferences with probabilities (Bayes) updated from the prior $[\pi, \rho, 1 - \pi - \rho]$ on the three states of the world and utility u^n . Let $w^n = (w_r^n, w_g^n, w_b^n)$ denote

an allocation of n 's wealth over the three state securities. Subject n 's total utility for such allocation equals

$$U^n(w^n) = \pi u^n(w_r^n) + \rho u^n(w_g^n) + (1 - \pi - \rho)u^n(w_b^n),$$

where we assume that the function u^n is twice differentiable, increasing and strictly concave. That is, each subject is risk averse.

Let $p = [p_r, p_g, p_b]$ denote the vector of state security prices, normalized so that $p_r + p_g + p_b = 1$. Subject n 's budget constraint equals:

$$pw^n \leq pw^{0,n},$$

where $w^{0,n}$ is the vector of subject n 's endowments of state securities. The first-order conditions for the optimal choice of state securities $w_s^{n,*}$, $s = r, g, b$, are standard:

$$\pi_s u^{n'}(w_s^{n,*}) = \lambda^{n,*} p_s,$$

where $\lambda^{n,*}$ is the Lagrange multiplier (corresponding to the budget constraint), and $\pi_r = \pi$, $\pi_g = \rho$, $\pi_b = 1 - \pi - \rho$.

The following are well-known implications of these first-order conditions.

- Because first-order conditions must be satisfied with equality, agents will always be marginal. In particular, they will adjust portfolio positions whenever prices changes.
- Since all agents face the same state probabilities and prices, and if, as is usually assumed, marginal utility is decreasing for all agents, then allocation of wealth to states will be perfectly rank-correlated across subjects. This means, if one were to rank states depending on the wealth implied by an agent's portfolio choice, any other agent's wealth ranking across states would be the same.

Neither implication will obtain under ambiguity.

Perfect rank correlation obviously does not imply perfect correlation. Perfect correlation only obtains if different subjects exhibit preferences that are members of particular families, such as quadratic preferences. In practice, perfect rank correlation does not even obtain; instead, portfolio weights are such that the typical investment (the statistical mode) reflects the relative per-capita supplies of the state securities and all subjects are spread out around that mode. We shall illustrate this in the empirical section of this paper.

Suppose next that subject n perceives ambiguity in the portfolio choice problem, and her attitudes towards risk and ambiguity are summarized in terms of α -maxmin preferences (see Ghirardato e.a. (2004)), with (state-independent) utility function u^n and the following set of priors:⁷

$$C = \{[\pi_r, \pi_g, \pi_b] \in \Delta^2 : \pi_r = \pi\}.$$

⁷ Δ^2 denotes the two-dimensional unit simplex in \mathbb{R}^3 .

As for an expected utility subject, the number π gets updated according to Bayes' law as information is received on the urn composition.⁸ In this case, subject n 's total utility for an allocation w^n equals

$$U^n(w^n) = \pi u^n(w_r^n) + \alpha_n \min_{\rho \in [0, 1-\pi]} \{\rho u^n(w_g^n) + (1-\pi-\rho)u^n(w_b^n)\} + (1-\alpha_n) \max_{\rho \in [0, 1-\pi]} \{\rho u^n(w_g^n) + (1-\pi-\rho)u^n(w_b^n)\},$$

where u^n is again assumed to be twice differentiable, increasing and strictly concave.

Intuitively, the coefficient α_n measures subject n 's aversion to the ambiguity described by the set C . When $\alpha_n = 1$ —that is, when subject n has the maxmin preferences of Gilboa and Schmeidler (1989)—the subject is extremely ambiguity averse. When $\alpha_n = 0.5$, the subject behaves like an expected utility maximizer with beliefs $[\pi, (1-\pi)/2, (1-\pi)/2]$, so that she looks neutral with respect to ambiguity.⁹

The first-order condition for the optimal quantity of state security r , $w_r^{n,*}$, is identical to that seen earlier:

$$\pi u^{n'}(w_r^{n,*}) = \lambda^{n,*} p_r.$$

The first-order conditions for the optimal choice of state securities g and b are more complicated. If $w_g^{n,*} > w_b^{n,*}$ (i.e., it is optimal for subject n to hold more state securities g than b), then the following conditions must be satisfied

$$\begin{aligned} (1-\alpha_n)(1-\pi) u^{n'}(w_g^{n,*}) &= \lambda^{n,*} p_g, \\ \alpha_n(1-\pi) u^{n'}(w_b^{n,*}) &= \lambda^{n,*} p_b. \end{aligned} \tag{1}$$

If on the other hand $w_b^{n,*} > w_g^{n,*}$ (i.e., it is optimal for subject n to hold less state securities g than b), then

$$\begin{aligned} \alpha_n(1-\pi) u^{n'}(w_g^{n,*}) &= \lambda^{n,*} p_g, \\ (1-\alpha_n)(1-\pi) u^{n'}(w_b^{n,*}) &= \lambda^{n,*} p_b. \end{aligned} \tag{2}$$

If $\alpha_n \in [0.5, 1)$, inspection indicates that the former conditions, i.e., (1), obtain if

$$\frac{p_g}{p_b} < \frac{1-\alpha_n}{\alpha_n}, \tag{3}$$

whereas (2) obtain if

$$\frac{p_g}{p_b} > \frac{\alpha_n}{1-\alpha_n}. \tag{4}$$

If $\alpha_n = 1$, neither set of conditions can obtain for strictly positive prices.

⁸If, for instance, the urn initially contained 18 balls, 6 of which were known to be red, then following an extraction of a red ball the subject uses $\pi = 5/17$. If a blue ball is extracted next, π is “updated” to $5/16$.

⁹This “similarity” of (1/2)-maxmin and expected utility preferences is not in general true, but it holds for the state space and set of priors we are using here. However, notice that dynamically an (1/2)-maxmin subject will behave differently from an expected utility maximizer, as the probabilities she attaches to states g and b will in general diverge from those obtained by Bayes' rule.

Outside of the bounds determined by (3) and (4), the subject chooses to hold equal quantities of the ambiguous state securities. In that case, subject n 's optimal holdings in both states, $w_s^{n,*}$, $s = g, b$, satisfy the following condition:

$$(1 - \pi) u^{n'}(w_s^{n,*}) = \lambda^{n,*}(p_g + p_b). \quad (5)$$

Notice that there is a region for the price ratio p_g/p_b where subject n is *not exposed to risk in the ambiguous states*. That is, conditional on one of the ambiguous states occurring, subject n always receives the same. For p_g/p_b above a certain level, subject n chooses to hold more of state security b than of g . Conversely, below another level, subject n holds more of g than of b . Figure 1 depicts these regions as a function of α_n . For $\alpha_n = 0.5$, subject n is always exposed to conditional risk unless $p_g = p_b$. When $\alpha_n = 1$, subject n is *never* exposed to conditional risk as long as prices are strictly positive.

The implications of the presence of ambiguity averse agents are fundamental.

- Ambiguity averse agents may not be marginal. Technically, this means that their first-order conditions for allocation of wealth across ambiguous states will not be satisfied with equality. When there are also ambiguity neutral agents, we obtain a situation in which some agents are marginal and others are not.
- Since ambiguity averse agents may want to hold portfolios that ensure equal wealth across ambiguous states, while ambiguity neutral agents will always choose portfolios that imply unequal wealth across states with different state price-probability ratios, perfect rank correlation of wealth across agents will no longer obtain.

Both implications are fundamentally different from those obtained under pure risk.

In the sequel, we shall assume that heterogeneity is such that some agents are indeed infra-marginal for ambiguous states. It is obvious how one would discern the presence of such agents: infra-marginal agents will hold portfolios such that their wealth in state g is equal to that in state b . When expressing their portfolio holdings in terms of state securities, this means that their holdings of state security g as a proportion of their holdings of state securities b and g is exactly 50%. In contrast, ambiguity-neutral agents will hold more or less than 50% in g , depending on the relative prices of the two state securities, which, in turn, are determined by the relative supplies of these securities.

We mentioned before that, in practice, one never observes perfect rank correlation of final wealth across states, even under pure risk. Instead, the mode of the distribution of portfolio holdings reflects the relative supplies of state securities. Under ambiguity, and when there are ambiguity averse agents who do not wish to be exposed to ambiguity, the distribution of portfolio holdings should change significantly. In particular, we should see a new mode, reflecting the desire of ambiguity averse subjects not to be exposed to ambiguity. When the supply of state securities is uneven, however, this implies that the ambiguity tolerant subjects need to compensate, by exposing themselves more to risk. As a result, the distribution of portfolio holdings may become heavily skewed, with ambiguity tolerant subjects over-

exposing themselves to risk (relative to the per-capita supply of aggregate risk), forcing the first mode to move away from the market portfolio weight. This is exactly what we observe in the experiments, as discussed later.

2.2 Equilibrium

Equilibrium is defined in traditional ways: the equilibrium price vector p^* is such that markets clear with demands equal to their optimum given p^* :

$$\frac{1}{N} \sum_{n=1}^N w_s^{n,*}(p^*) = \frac{1}{N} \sum_{n=1}^N w_s^{0,n}.$$

(The dependence of optimal demands $w_s^{n,*}$ on prices p^* is made explicit in this equation.) In the sequel, let W_s^0 denote the per-capita supply of state security s , i.e.,

$$W_s^0 = \frac{1}{N} \sum_{n=1}^N w_s^{0,n}.$$

In the case all agents have expected utility preferences, it is well known (see, e.g., Bossaerts e.a. (2005)) that because asset markets are complete, under the present assumptions there is a *representative agent with state-independent expected utility preferences*, who finds it optimal to hold the per-capita endowment of state securities at equilibrium prices. In particular, this finding implies that the equilibrium *state price probability ratios*, i.e., the ratios of the equilibrium prices over the beliefs (forming the state price density), will be inversely related to the social endowments of the state securities.

Why does this obtain? The answer is very simple: everyone is marginal, which, as mentioned before, implies perfect rank correlation of wealth. This also implies perfect rank correlation of individual wealth with the social endowment across states. Hence, one can construct a state-independent utility function of aggregate per-capita wealth that also satisfies first-order conditions (and features decreasing marginal utility). This utility function defines the preferences of the representative agent.

Note that the presence of heterogeneity in risk attitudes does not really impact *the structure* of equilibrium prices. We could have seen the same prices as the result of a market interaction of agents who all have the same risk attitude, as formally described by the utility function of the representative agent. Loosely put, a market with agents who do not care about ambiguity will not have (in equilibrium) any scope for heterogeneity in the agents' tastes.

Now assume, in contrast, that subjects are sensitive to the ambiguity in the problem, so that subjects' preferences and portfolio choices are as described in Section 2.1 (with common set of priors C). We begin by showing that if the agents have common ambiguity aversion coefficient, the result obtained above extends. That is, the heterogeneity of the agents —i.e., their different utility functions— does not really affect the structure of the equilibrium prices, when a positive price equilibrium exists.

If $\alpha_n = 1$ for all n —i.e., all agents are extremely ambiguity averse— then it is immediate to see that no equilibrium *with positive prices* exists (for any positive prices subjects will not want to be exposed to ambiguity; hence, as long as ambiguous securities are in unequal supply the market is not in equilibrium).¹⁰

If, on the other hand, $\alpha_n = \alpha \in [0.5, 1)$ for all n —i.e., all agents have common but not extreme ambiguity aversion— then it is simple to see that the equilibrium (exists with positive prices and) is identical to what obtains if all subjects maximize expected utility with respect to some common prior (which does not get updated as Bayes’ rule predicts, though). For instance, in the case in which $W_g^0 > W_b^0$, the equilibrium will be that obtained with expected utility subjects whose common beliefs are given by $[\pi, (1 - \alpha)(1 - \pi), \alpha(1 - \pi)]$ (and utilities given by the u^n ’s). A symmetric equilibrium is obtained in the $W_g^0 < W_b^0$ case. (There is a continuum of equilibria in the case in which $W_g^0 = W_b^0$. In that case many common priors exist that represent the subjects equilibrium beliefs.) Thus, we see that even in this case we can prove the existence of a representative agent, whose equilibrium beliefs coincide with the subjects’ implied equilibrium beliefs (as described above) and who finds it optimal to hold the social endowment of the state securities.¹¹ As a consequence, the result that the state price probability ratios (deflated using the agent’s implied equilibrium beliefs) are inversely related to the social endowments of the state securities extends to this case, proving the claim we made above.

But the available evidence suggests that homogeneity of perceptions and attitude to ambiguity is not a very plausible assumption. Subjects display significant differences in their perception of ambiguity, and in their response to such perception. Therefore, we now consider a “mixed” situation where subjects may display heterogeneity in such characteristics. As we shall see, such heterogeneity *will* affect the equilibrium prices, yielding in certain conditions a price structure that could not be observed in a representative agent equilibrium like the ones analysed thus far.

Suppose that the agents are of two types. Type-I agents are either ambiguity-insensitive expected utility maximizers, with prior beliefs given by $[\pi, \rho, 1 - \pi - \rho]$, or ambiguity-sensitive with α -maxmin preferences where $\alpha < 1$. Type-II agents are extremely ambiguity averse and have 1-maxmin preferences.

From Figure 1, for a given level of p_g/p_b , 1-maxmin subjects will never be exposed to conditional risk in the ambiguous states. One would therefore conjecture that only subjects with expected utility preferences or with $\alpha < 1$ (or with expected utility) will be marginal, and hence, determine the exact relative pricing of state securities g and b . For instance, if the total supply of state security b is inferior to that of g , then these subjects will push p_g below p_b sufficiently for them to be willing to accommodate the unequal supplies. Our previous analysis shows that, in this case, the α -maxmin agents would assign beliefs $[\pi, (1 - \alpha)(1 - \pi), \alpha(1 - \pi)]$ to the states.

¹⁰The positivity of prices is the reason behind the apparent discrepancy with the result of Epstein and Wang (1994) that an equilibrium exists (in a general model with 1-maxmin preferences, which encompasses the case we discuss here). In general, their equilibrium may have zero prices for some assets.

¹¹This result is not too surprising. See Epstein and Wang ((1994)), Appendix C and Dana (2004) for some more general representative agent theorems for the case of 1-maxmin preferences.

To determine the equilibrium pricing restrictions in more detail, order agents them so that agents $n = 1, \dots, N^*$ ($N^* < N$) are of Type I, i.e., have expected utility preferences or α -maxmin preferences with $\alpha < 1$, and agents $n > N^*$ are of Type II, i.e., have 1-maxmin preferences. Because everyone is marginal for state security r , the equilibrium condition for the market of the pure-risk state security r is standard:

$$\frac{1}{N} \sum_{n=1}^N w_r^{n,*}(p^*) = W_r^0.$$

(Again the dependence of optimal demands $w_r^{n,*}$ on prices p^* is made explicit in this equation.) In contrast, the equilibrium conditions for the two ambiguous markets are nonstandard:

$$\begin{aligned} \frac{1}{N^*} \sum_{n=1}^{N^*} w_g^{n,*}(p^*) &= \frac{N}{N^*} W_g^0 - W_a^*, \\ \frac{1}{N^*} \sum_{n=1}^{N^*} w_b^{n,*}(p^*) &= \frac{N}{N^*} W_b^0 - W_a^* \end{aligned} \quad (6)$$

where

$$W_a^* = \frac{1}{N^*} \sum_{n=N^*+1}^N w_g^{n,*}(p^*) = \frac{1}{N^*} \sum_{n=N^*+1}^N w_b^{n,*}(p^*)$$

(using: $w_g^{n,*} = w_b^{n,*}$ for $n > N^*$).¹²

These equations summarize the impact of heterogeneity in ambiguity aversion on prices. For state security r , *all* N agents determine the price; while for state securities b and g , only N^* agents determine the relative pricing. The relevant supply that determines prices of state securities b and g is not the total supply, but the total supply *minus* the holdings of the 1-maxmin agents, which are fixed across the two states.

To understand some of the consequences, imagine that $W_b^0 < W_g^0$, i.e., the per-capita endowment in state b is less than in state g . Under pure risk, equilibrium requires that the price-probability ratio of b has to higher than that of g , to refrain agents from buying state securities b .

Assume, for simplicity, that everyone's risk aversion is the same. Under ambiguity, and assuming that the $(N - N^* - 1)$ 1-maxmin agent would like to hold $W_b^0 < w_b^{n,*}(= w_g^{n,*}) < W_g^0$, the N^* ambiguity-tolerant agents need to absorb the *entire* supply imbalance across states b and g . They will do so only if the differential between the price-probability ratios of b and g is greater than under pure risk.

We consider two cases. First, imagine that $W_b^0 < W_r^0 < W_g^0$. Under pure risk, this would lead to a reverse ranking of the state price-probability ratios: the state price-probability ratio for b is largest; that of g is highest. Because of ambiguity, the differential between the state price-probability ratio of b and g increases. It is easy to see, however, that this will never lead to a change of the ranking of state price-probability ratios relative to the pure risk case. Call this the NRR case, for *No Rank Reversals*.

¹²To put this differently: subjects $n = 1, \dots, N^*$ are willing to hold the ambiguous securities in unequal quantities; they accept exposure to risk in the ambiguous states.

Second, consider the case where $W_b^0 < W_g^0 < W_r^0$. Again, the differential between the state price-probability ratio of b and g increases relative to the pure-risk case. However, now the state price-probability ratio for g may become larger than that of r , causing a change in the rankings. Whether it does depends on the amount of risk aversion and the number of 1-maxmin *vs.* ambiguity tolerant agents. We refer to this as the PRR case, for *Potential Rank Reversals*.

The above analysis relies on the assumption that risk aversion of all agents is the same and the total amount of ambiguous state securities the 1-maxmin agents buy. One can imagine situations where rank reversals may obtain even in the NRR case - e.g., if ambiguity tolerant agents are less risk averse, and 1-maxmin agents buy few of the ambiguous-state securities.

As such, the pricing implications of heterogeneity in ambiguity aversion are less obvious than the implications for allocations. One obtains “potential” rank reversals in certain situations, but even outside these situations, one *may* see rank reversals.

3 The Experiments

The experiments consisted each of a sequence of trading sessions, referred to as periods. At the beginning of a period, subjects were given a fresh supply of securities and cash. Markets opened and subjects were free to trade your securities. Subjects bought securities with cash and got cash if they sell securities. At the end of the period, after markets close, the securities expired, after paying dividends as specified below. These dividends, together with the cash balance, constituted subject period earnings. Period earnings were cumulative across periods. At the end of the experiment, the cumulative earnings were paid out, in addition to a standard sign-up reward.

In some experiments, we used a fictitious currency called francs, converted to dollars at a pre-announced rate; in other experiments, the accounting was done in real dollars. The results we are about to report does not depend on whether a fictitious currency or real dollars were used.

At the beginning of a period, subjects were endowed with 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. That is, stocks X, Y and Z were state securities.

In all experiments, subjects were also endowed with cash. In some experiments, part of this cash was considered a loan and subjects were asked to pay back the loan out of end-of-period proceeds. In other experiments, subjects started with a negative endowment of bonds, effectively meaning that they were given a loan. In the sequel, we report loans as negative endowments of bonds.

Subjects needed cash to buy stocks and bonds. Subjects could short-sell stocks and bonds, but were not allowed to take positions that may cause losses of more than \$2. In earlier experiments, we imposed this limit *ex post*, by barring subjects with more than \$2 losses from trading in future periods; in later experiments, we 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, the pending order was not submitted.

Trading took place over an electronic market organized as a continuous open-book system (a type of double auction where infra-marginal orders remain displayed until either executed or canceled). Three different interfaces were used, which all generated qualitatively the same results: (i) an interface where orders (both quantities and prices) had to be entered manually, *Marketscape*;¹³ (ii) an interface where market orders (orders that executed immediately at the best available price) could be entered by clicking only, *eTradeLab*;¹⁴ (iii) an interface where all orders are submitted by point-and-click, *jMarkets*.¹⁵

Here is how states were 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 drew one ball from this urn. Imagine that we drew an X ball. This will determine the payoff on the stocks: for each unit of stock X subjects were holding at the end of the period, they received \$0.50. You did not receive anything for holdings of stocks Y and Z. We then threw away this ball (drawing without replacement). As a consequence, we drew the state for period 2 from an urn with the following composition: 6 X balls, 7 Y balls, and 6 Z balls. If we drew a Z ball in the second period, then the state for the third period would be drawn from an urn with the following composition: 6 X balls, 7 Y balls, and 5 Z balls. Etc.

We drew the state at the beginning of each period. This meant that the state did not depend on what subjects did during the period. Nobody was told what the state was until the end of the period.

A typical experiment lasted for 2 1/2 hours. Subjects earned between \$0 and \$125, with an average of approximately \$50.

Experiments started with two practice periods. In some experiments, subjects were paid for their end-of-period holdings in the practice periods; in others, they were not, which also meant that they could explore more – mistakes had no earnings consequences. We did not see any material differences between experiments where subjects were paid during practice periods and those where they were not.

In all experiments, the total endowment of X securities was smallest, while that of Y securities was largest. Subjects were not given this information.

Table 1 displays all the treatments we ran. Each experiment was either of the NRR or PRR type. In NRR

¹³Marketscape was developed in Charles Plott's lab.

¹⁴eTradeLab was developed by Tihomir Asparouhov.

¹⁵jMarkets is open-source software, developed at Caltech, and available at <http://jmarkets.ssel.caltech.edu>.

experiments, states X and Y were ambiguous in that we did not tell subjects what proportion of balls in the urn were X or Y. Table 1 displays the actual composition of the urn. The analysis of the previous section suggested that state price-probability ratio rankings in the NRR treatment would (likely) be as under pure risk: the state price-probability ratio of X would be highest; that of Y would be lowest.

In the PRR type, states Y and Z were ambiguous. Reversals in rankings of state price-probability ratios may now be possible, but, as the analysis in the previous section indicated, such reversals are not necessary.

Three sets of experiments were run for each treatment: A, B, C. In the cases of A and B, the ambiguity experiments were replicated as pure-risk experiments. These pure-risk replications were identical in that the same urn was used, as well as *the same draws*. As indicated in Table 2, approximately the same number of subjects were used, but their identity was different (to avoid any potential prior knowledge of the draws, the pure-risk replications were ran more than one year later).

Subjects were differentiated in terms of initial allocations. See Table 3. To interpret the distributions of final-period holdings, it is important to determine what these initial allocations, together with the number of subjects of each type (see Table 2) implies in terms of total wealth available across states. Table 4 provides details. It shows the total wealth in state Y as a fraction of the total wealth in the two ambiguous states (X and Y in the NRR treatment; Y and Z in the PRR treatment). We will refer to these as the *ambiguous market portfolio weight of Y*.

Highly ambiguity averse subjects (1-maxmin subjects) are expected to invest in a portfolio so that their wealth in state Y as a fraction of total wealth over the ambiguous states equals 0.5. This means that they will not be exposed to ambiguity: they receive the same payment in the two ambiguous states. The ambiguity tolerant subjects are expected to compensate: the modal subject should weigh state Y (as a proportion of total wealth in the two ambiguous states) than the market portfolio. Under pure risk, the modal investor should invest to a pattern of risk across ambiguous states that corresponds to that of the ambiguous market portfolio.

Note that the ambiguous market portfolio weight of Y is much higher in the NRR treatment than in the PRR treatment. In particular, the ambiguous market portfolio weight in the PRR treatment, at 0.62, is much closer to 0.5, the weight that 1-maxmin subjects were expected to implement.

The reader can expect the web site for the experiments:

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

As of writing of this draft, web pages reflect the last experiment we ran, which is the pure-risk replication of Treatment (NRR, B).

4 Empirical Results

The state price-probability ratios (ratios of state prices over probabilities, or the state price density) together with cross-sectional patterns of holdings of ambiguous securities and of exposure to risk provide insight in the nature of ambiguity attitudes among subjects, the potential presence of heterogeneity in this attitude, and the impact on pricing. We first discuss state price-probability ratios. Subsequently, we study end-of-period holdings of securities and their implication for the distribution of wealth across states. Finally, we use end-of-period holdings to estimate correlation between risk and ambiguity aversion.

Recall that here states and securities are called X, Y and Z, with the risky state being Z in some experiments and X in the others, and that, in the notation of the previous section, $W_X^0 < W_Z^0 < W_Y^0$ in each experiment.

4.1 State Price-Probability Ratios

To determine the state price-probability ratios, the state prices were divided by the probability of the corresponding states, assuming initially a uniform prior (over the ambiguous states), suitably updated using Bayes' law. Other choices are possible; we will elaborate in the Discussion section.

We display state pricing results in the form of empirical distribution functions (ECDFs) because these provide unbiased estimates of the probabilities that state price probability ratios exceed any given level, unaffected by obvious time series considerations such as autocorrelation and conditional heteroscedasticity. That is, they provide unbiased answers to questions of the type “is there a higher chance that the state price probability ratio of X is above 1 than the state price probability ratio of Z?” Moreover, by the Glivenko-Cantelli theorem, empirical distribution functions converge uniformly, so that statements of first-order stochastic dominance can meaningfully be made (e.g., “the probability that the state price probability ratio of Y exceeds any level is greater than that for X”).

Absent precise knowledge of subjects' actual attitudes towards risk and ambiguity, only the ordinal properties of state price probability ratios can be investigated. Moreover, markets go through lengthy adjustments even in situations as simple as the present ones. This means that many if not most transactions take place before markets have settled, if they settled at all. In fact, it is not clear when markets settle, because bid-ask spreads in our experiments were generally large, unlike in experiments without ambiguity —an obvious sign that some subjects did exhibit aversion against ambiguity.¹⁶ Because of the sizeable bid-ask spreads, price volatility remained high throughout the experiment.

We first discuss the results in the A and B versions of our experiments, because it allows us to contrast pricing in ambiguous situations with that in analogous pure-risk situations.

First, consider treatment NRR. Remember that in this treatment, X and Y were ambiguous, so that ranking

¹⁶Under ambiguity aversion, there is a region of prices for which a subject would not trade. See Figure 1. This region corresponds to the interior of the bid-ask spread.

violations of state price-probability ratios are unlikely once ambiguity is introduced. Figure 2 displays the ECDFs of the state price-probability ratios in the NRR pure-risk experiments A and B and their ambiguous counterparts. Under pure risk, state X should be most expensive (relative to its probability), while state Y should be the cheapest. This is confirmed in the data. The top panels of Figure 2 display the results for the pure-risk NRR experiments (left: A; right: B). The ECDFs of the state price-probability ratio of Y indeed first-order stochastically dominate that of Z, and the ECDFs of Z first-order stochastically dominate that of X. Below them are the ECDFs of the corresponding ambiguity experiments. Notice that the ranking of ECDFs (in terms of first-order stochastic dominance) is preserved.

Next, consider treatment PRR. In this treatment, Y and Z were ambiguous states. Ambiguity aversion may push the price of Z (relative to its probability) to become higher than that of X, upsetting the ranking property that obtains under pure risk. The top panels in Figure 3 show the ranking of the ECDFs in the pure-risk version of these experiments. In experiment A (left side), ranking is as predicted in the theory. In experiment B (right side), the ECDF of Z is actually to the right of that of Y for most of the range of state price-probability ratios, contradicting the theory. Such violations have been observed before [see Bossaerts and Plott ((2004))] and generally occur when one state is drawn so often relative to another one that subjects believe that the former won't occur anymore while the latter is "overdue" (relative to rational updating). In the (PRR, B) experiment, such a situation obtained: state Z was drawn 4 (out of 6) periods, while state Y was never drawn. As a result, Z became very cheap, while Y became more expensive.

The bottom panels of Figure 3 display the ECDFs for the corresponding ambiguity experiments. In experiment A (left graph), we obtain a ranking violation: as predicted, Z becomes expensive so often that its ECDF is actually for a large part to the right of the ECDF of X. Whereas it is almost certain to observe a price of X of 1.5 (relative to its probability), there is only about 70% for this to occur with the price of Z. We do not observe a ranking reversal in experiment B. However, relative to its pure-risk counterpart, the price of Z has indeed shifted upward: the ECDF of Y now first-order stochastically dominates that of Z, while it did not in the pure-risk version of this experiment. As such, the impact of the presence of ambiguity averse subjects is as expected.

We ran two more experiments under the ambiguity treatment only (C). Figure 4 displays the ECDFs of the state price-probability ratios in these experiments. The left panel shows the results for the NRR version: we expected ranking of ECDFs as under pure risk, and indeed we obtain this. The right panel shows the ECDFs for the PRR version: rank reversals becomes possible. While in the pure-risk case, the ECDF of Y should first-order stochastically dominate that of Z, it no longer does so under ambiguity. Notice however that, in contrast to the PRR cases plotted in Figure 3, the prices of Y and Z are now closer together. In theory, this can obtain if ambiguity tolerant agents (who need to absorb the difference in total supplies of state securities Y and Z) are closer to risk neutral than the population as a whole. Later on, we shall document that ambiguity tolerant subjects indeed tend to be less averse than ambiguity averse subjects.

As such, the pricing effects from the introduction of ambiguity are as expected. Rank changes in state price-probability ratios obtain only in the PRR treatment, when the security in shortest supply pertains to a state (X) that is pure risk, while the other states are ambiguous.

4.2 Final Holdings

Prices have given us some indication that there is heterogeneity in attitudes towards ambiguity. End-of-period holdings provide strong confirmation.

The effects of ambiguity aversion on portfolio choice can best be appreciated by comparing the final holdings in experiments with ambiguity to those in exact replications without ambiguity. Under pure risk, everyone is expected to allocate as much to Y as is available per-capita. This prediction under pure risk requires that everyone fully optimizes, that all preferences are smooth and that all beliefs are equal. In practice, only the modal person will follow these precise prescriptions, as we shall show. Under ambiguity, we expect to see peculiar distributions of end-of-period exposures. In particular, the ambiguity averse subjects will trade towards a portfolio that is not exposed to ambiguity. This requires that 50% of wealth in the two ambiguous states be allocated to each of them. The ambiguity tolerant subjects need to compensate, because there is more aggregate wealth in one of the ambiguous states (Y). That is, the ambiguous market portfolio weight of Y is higher than 50%, and ambiguity tolerant subjects will have to trade to a portfolio with more exposure to Y than the ambiguous market portfolio. As a result, the mode of the distribution of the exposure of Y as a percentage of total exposure to ambiguous states will be no longer at the ambiguous market portfolio weight, but at 50%. To compensate, the tails will be longer - especially the right tail.

Figure 5 compares the histograms of final holdings in the NRR treatment. The top panels provide the results for the pure-risk case (left: experiment A; right: experiment B). The bottom panels provide the results for the corresponding cases with ambiguity. Figure 6 displays the results for the PRR treatment. In all pure-risk cases, the distribution of exposure to Y is symmetric around a mode that equals the ambiguous market portfolio weight (i.e., the proportion assigned to Y in the market portfolio of Y and the other state that is ambiguous under the ambiguity treatment). This ambiguous market portfolio weight is higher in the NRR treatment (0.98; see Table 4) than in the PRR treatment (0.63), and this is reflected in the pure-risk distributions of Figures 5 and 6 as well.

The theory predicts a sharply different distribution of final holdings under ambiguity, and Figures 5 and 6 (bottom panels) confirm this. The mode of the distributions have shifted, to 50%, reflecting the actions of ambiguity averse subjects. They have also become more heavy-tailed, mostly in the right tail, reflecting the compensating choices of ambiguity tolerant subjects. This reveals strong evidence of the presence of heterogeneity in attitudes towards ambiguity.

Experiment C was only ran under the ambiguity treatment. Figure 7 displays the distribution of final portfolio

holdings for the NRR and PRR versions of experiment C. Again, we confirm the peculiar mode (at 50%) of these distributions, and their heavy (right) tails.

4.3 Correlation Between Risk And Ambiguity Aversion

As a by-product of our experiments, we can investigate to what extent subjects' attitudes towards risk and towards ambiguity correlate. In theory, risk aversion is determined by the concavity of an agent's utility function, while the parameter α controls ambiguity aversion.

To measure a subject's aversion to ambiguity, we record the imbalance in wealth across ambiguous states that are implied by his/her end-of-period portfolio holdings. To measure risk aversion, we compute the maximum difference in wealth across all states that end-of-period portfolio holdings entail.

Figure 8 displays the results for all subjects and all periods of all the experiments involving ambiguity. There is a pronounced tendency for more risk tolerant subjects (whose portfolios imply a larger range in wealth across all states) to be more ambiguity tolerant as well (their wealth across ambiguous states is also more unequal).

Such correlation between ambiguity and risk aversion is to be taken into account in theoretical analyses of asset pricing. It suggests, among other things, a novel explanation of the value effect, i.e., the tendency for stock in companies with high book value relative to market value to earn higher returns, and hence, to carry a higher risk premium. If low (book) value can be associated with ambiguity, then low-value securities are held mostly by ambiguity tolerant agents, who are also more risk tolerant, and hence, require a lower risk premium – implying that the require average return on low-value securities is lower. In contrast, high (book) value securities are held by everyone, including the more risk averse agents. With the right parametrization, this leads to a higher risk premium on the high-value securities.

The correlation between ambiguity and risk aversion is to be taken into account in theoretical analysis of potential effects of regulation on securities pricing and holding (Easley and O'Hara (2005)).

Our finding of positive correlation is foreshadowed in a recent study of the choices of patients with lesions in the orbito-frontal cortex and normal subjects [Hsu e.a. (2005)]: such patients tend not only to be more risk tolerant (as inferred from choices among risky securities) but also more ambiguity tolerant (as inferred from their optimism in situations where probabilities are unknown); normal subjects are more risk averse and tend to be ambiguity averse.

5 Discussion

There are two issues that deserve more discussion. The first one concerns our deflating prices of ambiguous states with *uniform* probabilities, suitably updated based on draws in past periods. Other choices are possible and we would like to discuss the effects here. The second issue concerns alternative explanations of the pricing and holding patterns.

We want to elaborate, in particular, on whether we really need heterogeneity (in attitudes towards ambiguity) and whether simple Savage-type expected utility would not be able to provide an alternative explanation.

First, we deflated state prices with probabilities based on the “obvious” uniform prior (over the ambiguous states) and updating according to Bayes’ law. Effectively, we deflated state prices using the beliefs of a Savage expected utility maximizer with uniform (conditional) prior beliefs over the ambiguous states. This also means that we assumed that the ambiguity tolerant subjects (who set the relative prices of the ambiguous states) are in fact Savage agents. One wonders to what extent our findings are sensitive to alternative assumptions about these probabilities. A distinct deflation would be to use a *fixed* uniform probability for the ambiguous states. This corresponds to the beliefs of an α -maxmin agent with $\alpha = 0.5$. Inspection of the results¹⁷ shows that state price-probability ratio ranking reversals continue to be absent in the NRR treatment. In the PRR treatment, the ECDFs of the price-probability ratios of the two ambiguous states (Y and Z) fail to show clear stochastic dominance; in fact, they are very similar, suggesting that the relative pricing of the ambiguous states reflects less risk aversion than the relative pricing of the ambiguous *vs.* risky states. The fact that ambiguity tolerant are more risk tolerant may be explain this effect: the ambiguity tolerant subjects set relative prices of ambiguous states, and since they are more risk tolerant as well, they require less of a risk premium, which means that the discrepancy of the prices of the two ambiguous states will be less.

Second, are there other explanations for the observed prices and choices? In particular, do we really need heterogeneity, or even ambiguity aversion?

We don’t need heterogeneity if we only are to explain prices. One can always pick the right alpha and explain pricing patterns. It is sufficient to find an alpha and, hence, corresponding beliefs, such that the ECDFs of the state price-probability ratios always display the correct ranking, as if there were a single representative agent with that alpha. Alpha would need to be different across experiments, however.

But the final holdings patterns cannot be explained without appeal to heterogeneity. It is clear that there always is a large group of subjects who choose not to be exposed to ambiguity. These are ambiguity averse subjects. There must be another group whose choices accommodate the unequal supply of the ambiguous state securities. These subjects must be less ambiguity averse. As a result, there must be heterogeneity.

Do we really need ambiguity aversion or could we explain prices and choices on the basis of a representative agent with Savage expected utility? Again, it is possible to explain pricing patterns by picking the right a priori beliefs. For the NRR treatment, one could simply choose the uniform prior: this is the type of belief we used to deflate state prices and the resulting ECDFs did not display any anomaly (see Figure 2). A priori beliefs will have to be different, however, for the PRR treatment, because simply picking the uniform prior generates rank reversals, as Figure 3 demonstrates. The Savage expected utility explanation does lose credibility, however, if beliefs change with the distribution of wealth across ambiguous and pure-risk states (which is the only material difference between the NRR and PRR treatments).

¹⁷ECDFs for the fixed uniform deflator can be obtained from the authors upon request.

In addition, the Savage expected utility explanation cannot readily accommodate the presence of a large number of subjects who choose to not be exposed to the risk across ambiguous states. That is, it would be improbable coincidence to always observe a large group of subjects who allocate exactly 50% to state Y as a proportion to total allocations in the ambiguous states.

6 Concluding remark

Our theory does not make sharp predictions about prices. It makes sharper predictions about allocations, in particular, the presence of agents who choose not to be exposed to ambiguity and the compensation that others provide. The empirical results also display this difference: the holdings patterns are uniform across all experiments; prices are not.

Ours is, therefore, a nice example where asset pricing theory makes more successful predictions about portfolio choices than about prices.

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 (2005): “Prices and Portfolio Choices in Financial Markets: Theory and Experiment.” Caltech working paper.
- Cagetti, M., Hansen, L., T. Sargent and N. Williams (2002): “Robustness and Pricing with Uncertain Growth.” *Review of Financial Studies* 15, 363-404.
- Chapman, D. and V. Polkovnichenko (2005): “Heterogeneity in Preferences and Asset Market Outcomes,” Carlson School of Management, University of Minnesota, working paper.
- 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.
- 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.
- 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.
- 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.

Table 1: Initial Composition Of Urn

Experiment Type	Experiment	Urn ^a		
		X	Y	Z
NRR ^b	A	6	3	9
	B	9	3	6
	C	6	3	9
PRR ^c	A	6	6	6
	B	6	6	6
	C	6	6	6

^aAmbiguous states in red.

^bNo reversals in ranking of state prices relative to total state wealth expected even in the presence of ambiguity (states X and Y are ambiguous).

^cReversals in ranking of state prices relative to total state wealth are possible in the presence of ambiguity (states Y and Z are ambiguous).

Table 2: Number of Subjects of Each Type

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

^aNo reversals in ranking of state prices relative to total state wealth expected even in the presence of ambiguity (states X and Y are ambiguous).

^bNumber of subjects of Type I and II, respectively.

^cReversals in ranking of state prices relative to total state wealth are possible in the presence of ambiguity (states Y and Z are ambiguous).

Table 3: Experiment Parameters

Experiment Type	Subject Type	Endowments				Initial	Loss
		X	Y	Z	Notes	Cash	Limit
NRR ^a	I	0	28	0	0	\$2	\$2
	II	0	12	20	-5	\$1	\$2
PRR ^b	I	3	12	0	0	\$1.2	\$2
	II	0	4	9	-4	\$2	\$2

^aNo reversals in ranking of state prices relative to total state wealth expected even in the presence of ambiguity (states X and Y are ambiguous).

^bReversals in ranking of state prices relative to total state wealth are possible in the presence of ambiguity (states Y and Z are ambiguous).

Table 4: Equilibrium Allocations Of Ambiguous State Securities

Experiment Type	$\frac{\#Y}{\#X+\#Y}$ (^a)	$\frac{\#Y}{\#Y+\#Z}$ (^b)
NRR ^c	0.98	-
PRR ^d	-	0.63

^aNumber of units of security Y as a fraction of total number of securities that pay in ambiguous states (approximate; actual number depends on number of subjects of each type).

^bNumber of units of security Y as a fraction of total number of securities that pay in ambiguous states (approximate; actual number depends on number of subjects of each type).

^cNo reversals in ranking of state prices relative to total state wealth expected even in the presence of ambiguity (states X and Y are ambiguous).

^dReversals in ranking of state prices relative to total state wealth are possible in the presence of ambiguity (states Y and Z are ambiguous).




Figure 1: Ratio of the prices of security g over b above which an α -maxmin agent sells security g (top line) or below which (s)he buys (bottom line). α is on the X axis; $\alpha = 0.5$ corresponds to ambiguity neutrality; $\alpha = 1$ corresponds to extreme ambiguity aversion.

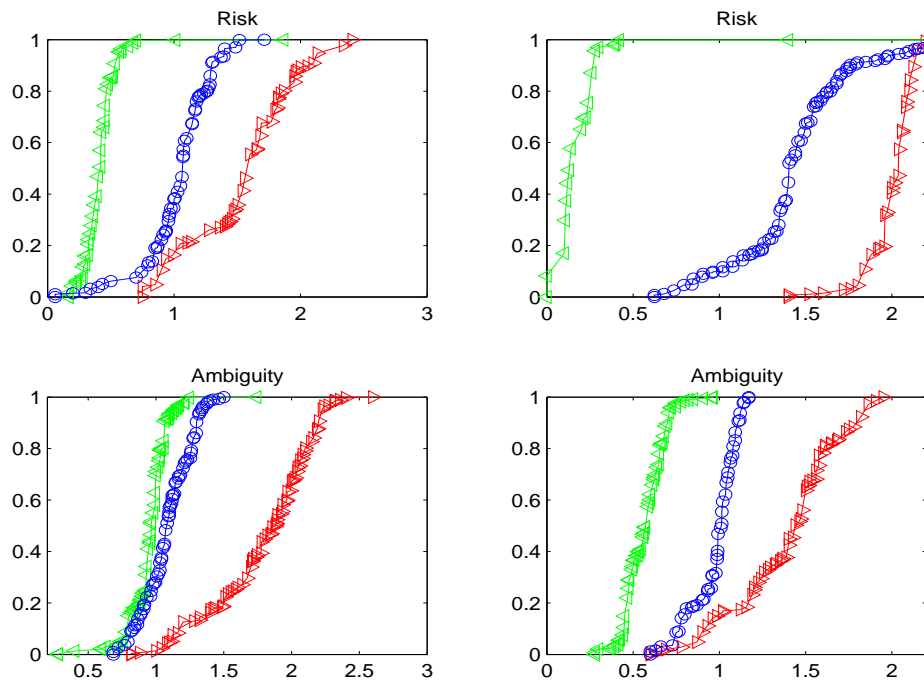


Figure 2: Empirical Distribution Functions (ECDFs) of state price-probability ratios, NRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.

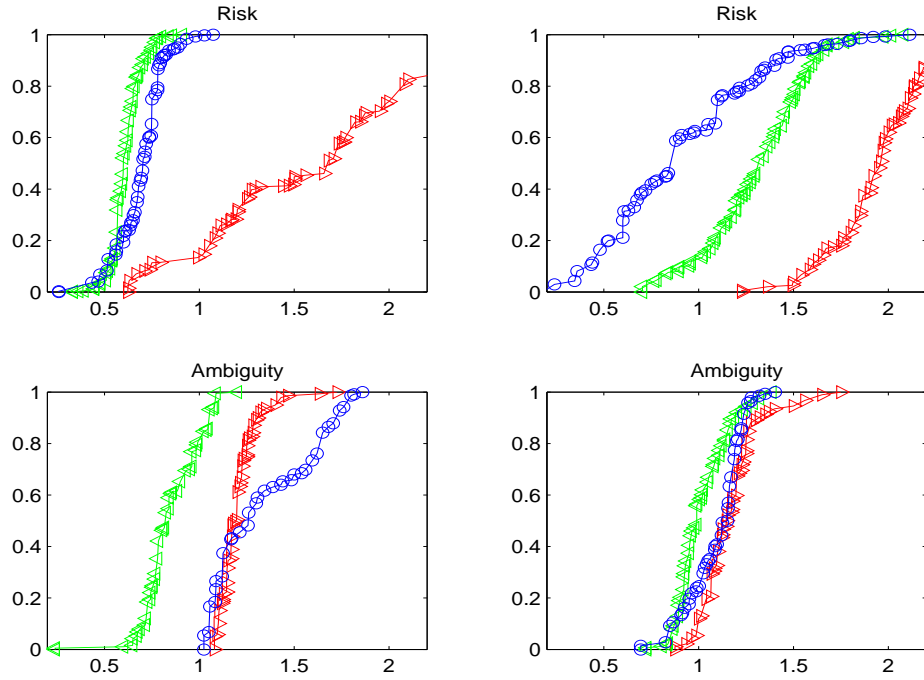


Figure 3: Empirical Distribution Functions (ECDFs) of state price-probability ratios, PRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.




Figure 4: Empirical Distribution Functions (ECDFs) of state price-probability ratios, ambiguity treatment (experiments C only; left: NRR treatment; right: PRR treatment).




Figure 5: Histograms of final wealth in state Y as a proportion of final wealth in states X and Y (the ambiguous states in the ambiguity treatment), NRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.




Figure 6: Histograms of final wealth in state Y as a proportion of final wealth in states Y and Z (the ambiguous states in the ambiguity treatment), NRR treatment. Top panels: pure-risk treatment (left: A; right: B); bottom panels: corresponding ambiguity treatment.




Figure 7: Histograms of final wealth in state Y as a proportion of final wealth in the two ambiguous states (X and Y in NRR treatment, left panel; Y and Z in PRR treatment; right panel).




Figure 8: Plot of difference from 0.5 of wealth allocated to state Y as a proportion of final wealth in the two ambiguous states (X and Y in NRR treatment; Y and Z in PRR treatment), against range of wealth allocated across all states; all periods in all experiments with ambiguous states.