Ambiguity and Asset Prices:
An Experimental Perspective

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

Preliminary and incomplete: Please do not circulate

[‡]Bossaerts: Caltech and CEPR; Guarnaschelli: McKinsey; Ghirardato: Università di Torino and University of Iowa; Zame: UCLA.

Ambiguity and Asset Prices:
An Experimental Perspective

Abstract

The violations of expected utility axioms displayed in the Ellsberg paradox have recently been attributed to ambiguity aversion. In this paper, we study the impact of ambiguity aversion on equilibrium asset pricing and portfolio holdings in competitive financial markets. We pay particular attention to potential heterogeneity, because a significant minority usually does not violate expected utility axioms. Our analysis is carried out in the context of state securities, some of which pay in states for which probabilities are unknown (the *ambiguous states*) but others pay in states for which probabilities are known (the *risky states*). Heterogeneity in ambiguity aversion leads to a wider range of state price probability ratios (state prices divided by probabilities, also known as state price density). If the ambiguous securities are not all in low or high supply, heterogeneity in ambiguity aversion could merely be misinterpreted as higher risk aversion. Otherwise it potentially generates violations to the ranking of state price probability ratios typical under expected utility, as if the representative agent held state-dependent utility. Experiments confirm the predicted impact. Heterogeneity in ambiguity aversion is further evident in subjects' end-of-period holdings. These holdings also reflect positive correlation between risk and ambiguity aversion. The latter suggests an explanation of the value effect, if value stock can be labeled 'pure risk' securities and growth stock 'ambiguous' securities.

Ambiguity and Asset Prices:
An Experimental Perspective[†]

1 Introduction

It has been suggested that ambiguity aversion explains the violations of the axioms of expected utility revealed in the Ellsberg paradox [(Ellsberg (1961))].[1] The expected utility model has therefore been generalized to allow agents to be averse to ambiguity, for instance by allowing agents to entertain multiple priors and to choose the actions which maximize the least expected utility over the possible priors (the so-called "maxmin expected utility" model of (Gilboa and Schmeidler (1989))).

In discussions about the Ellsberg paradox and ambiguity aversion, however, it is often overlooked that the data clearly show that not everybody violates the axioms of expected utility theory, and more generally that not everybody is ambiguity averse. That is, there is heterogeneity with respect to attitudes towards ambiguity —as could be represented, say, by an " α-maxmin" model in which agents weigh the least and highest expected utility with weights α and $(1-\alpha)$ (Ghirardato e.a. (2004)). The question we address here is: given such heterogeneity, what are the likely outcomes in a market setting where securities are traded whose payoffs have varying levels of ambiguity? What are the effects on equilibrium prices and allocations? Are these predictions upheld in laboratory versions of competitive financial markets?

Absent ambiguity, i.e., in the presence of pure risk, heterogeneity in aversion towards risk affects prices and allocations in ways that are well understood. First of all, if there is aggregate risk and nobody is risk neutral, *everybody* remains exposed to risk, no matter how risk averse. That is, all agents remain marginal. Second, it is not the (arithmetic) average risk aversion that is reflected in equilibrium prices, but the *harmonic mean risk aversion.* Because of the well known relationship between the arithmetic and harmonic mean, this implies that the more risk tolerant agents as a group have a disproportionately larger effect on pricing than do the more risk averse agents. Third, if markets are complete, the ratio of equilibrium state prices over state probabilities will be ranked inversely to the aggregate wealth across states. We will refer to these ratios as *state price probability ratios.*[2]

[†]Financial support from the U.S. National Science Foundation (grant SES-0079374), the R.G. Jenkins Family Fund and the Italian MIUR are gratefully acknowledged.

[1]The paradox that is closer to our experimental set-up is the following: An well-mixed urn contains 90 balls, 30 of which are red, while the others are green or blue, in an unspecified proportion. Would you prefer to bet on the extraction of a red ball or of a blue ball? Would you prefer to bet on the extraction of a red or green ball, or of a blue or green ball? Many subjects prefer "red" in the first choice and "blue or green" in the second.

[2]The mathematical finance literature refers to the state price probability ratios as *state price density.*

As we will demonstrate in this paper, the situation is substantially more complex when there is ambiguity and heterogeneity in attitudes towards ambiguity. Primarily, this is because the more ambiguity averse chose *not to be exposed to ambiguity*, and therefore, do not influence the relative prices of the ambiguous securities. They do have an effect on the pricing of ambiguous securities relative to pure-risk securities, further complicating matters.

Our analysis takes place in a simple static world with state securities. Agents do not know the probabilities of some states and we refer to the corresponding state securities as the *ambiguous securities*. When the probability of a state is known, the corresponding state security is merely *risky*. We study the effect on equilibrium pricing and the cross-section of portfolio holdings when we vary the relative supplies of the state securities.[3]

When there is heterogeneity in ambiguity aversion, the first thing to note is that the range and variance of equilibrium state price probability ratios is increased.[4] To see why this is, it is important to remember that, for most price configurations, highly ambiguity averse agents prefer to hold ambiguous securities in equal quantities. They leave the less ambiguity averse and the ambiguity neutral (including expected utility maximizers) to absorb the entire imbalance in relative supplies of ambiguous securities. The imbalance to be accommodated is far greater than if the ambiguity averse subjects had taken on their share, as they would in the absence of ambiguity (in which case *all* agents remain at least minimally exposed to risk, as already mentioned before). The less ambiguity averse subjects will absorb the supply imbalances in ambiguous securities only if compensated appropriately. Whence the increase in the spread (range; variance) of state price probability ratios.

Second, if ambiguous securities are in high supply, the more ambiguity tolerant agents may demand so much compensation for holding all the ambiguous securities that the simple ranking relationship between state price probability ratios and aggregate wealth is upset. When markets are complete and there is only pure risk, state price probabilities and aggregate wealth are inversely ranked, as mentioned before. When there are ambiguous states, the state price probability ratios of some of the ambiguous states may rise above those for pure-risk states, even when aggregate wealth in the latter states is lower. In other words, violations of the simple principle that states with lower aggregate wealth will be more expensive (per unit probability) are to be expected. An analogous possibility emerges when all the ambiguous securities are in lowest supply.

We ran a number of experiments with competitive financial markets where ambiguous securities traded alongside risky ones, in order to determine whether the theoretical effects of heterogeneity are borne out in the data. By and large, they are. Most importantly, we find violations of standard state price probability ratio rankings where the theory predicts they may occur; we don't find them when the theory rules them out. Furthermore, we confirm the

[3]Our setting contrasts with that of (Epstein and Miao (2003)). The latter studies asset pricing effects in two-agent economies where the two agents are equally ambiguity averse, but do not agree on what the ambiguous states are. In our setting, agents agree on the nature of the states, but exhibit differing levels of ambiguity aversion.

[4]The state price probability ratio for ambiguous states is defined as the ratio of the state price over a probability that reflects uniform initial priors over the ambiguous states, updated to reflect the history of states drawn in the past.

heterogeneity in aversion to ambiguity in the holdings, with some subjects holding ambiguous securities in approximately equal quantities so as not to be exposed to ambiguity, and others taking on unbalanced positions. At the same time, the holdings reveal positive correlation between risk and ambiguity tolerance.

This paper adds to an emerging literature that studies the effect of non-standard preferences on prices and choices in competitive markets through experiments. (Gneezy e.a. (2003)) analyzes the impact of myopic loss aversion on pricing, but, unlike this paper, assumes homogeneous preferences. (Kluger and Wyatt (2004)) studies the impact of well-known biases in updating on pricing in experimental markets, but, unlike this paper, provides no theoretical framework within which to understand how heterogeneity explains the experimental results.

The results of our experiments should also be contrasted with those emerging in experiments under pure risk. In that case, (Bossaerts e.a. (2003)) documents that there is substantial heterogeneity in subjects' demands. Nevertheless, demands are correlated in one important respect, namely, rationality. That is, demands are the sum of a common element predicted by rational choice theory under risk aversion, and an uncorrelated noise term. The correlation is small, but in large experiments, it is sufficient to induce pricing as if all choices had been made by fully optimizing and rational agents. Under ambiguity, however, theory predicts that heterogeneity may affect equilibrium pricing in certain circumstances. The experimental data confirm the predicted biases. Hence, under ambiguity, heterogeneity is not eliminated in the aggregate.

This finding has important implications for attempts to model pricing as if there were a homogeneous aggregate investor. Recently, several attempts have been made to explain the equity premium puzzle (high average returns on equity and low average riskfree rate) by appealing to ambiguity (or, equivalently, Knightean or model uncertainty) and assuming that there is a representative agent with simple state-independent utility and aversion to ambiguity (or uncertainty). The literature includes, but is not limited to (Epstein and Wang (1994), Uppal and Wang (2003), Cagetti e.a. (2002), Maenhout (2000)). Our experimental results indicate that the natural heterogeneity in attitudes towards ambiguity even in relatively small groups induces pricing that cannot be modeled simply as if there is an ambiguity averse aggregate investor with state-independent preferences.

Our findings suggest an explanation for the well-documented *value effect* in equity pricing. It has been documented extensively [see, e.g., (Fama and French (1993))] that the average return on growth stock has been small relative to that of growth stock, after accounting for differences in risk. To the extent that growth stock can be associated with ambiguity and value stock represent pure risk, natural segmentation of financial markets under heterogeneity in attitudes towards ambiguity would ensue. Growth stock would be priced only by ambiguity neutral investors, and hence, reflects only their risk aversion. If our experiments are any guidance for what is to be expected in the population at large, our finding that ambiguity neutral subjects are also less risk averse would imply a low risk premium on growth stock, and hence, low average returns. In contrast, the pure-risk value stocks are priced by all investors, so they carry a higher risk premium, and hence, higher average returns.[5]

The remainder of the paper is organized as follows. In the next section, we present the experimental setup. The framework of risk and ambiguity is analyzed theoretically in Section 3, generating asset pricing and portfolio holding predictions. Section 4 analyzes the data in view of these predictions. Section 5 provides a final remark concerning the relationship with pricing under expected utility but heterogeneous beliefs.

2 The Experiments

The experiments are organized as follows. Three state securities can be traded, referred to as securities X, Y and Z, corresponding to the labeling of the states. A fourth asset, called Notes, is riskfree, and, unlike the risky securities, can be sold short, up to eight units. At the beginning of each period, subjects are endowed with a certain number of the state securities and cash (purchases are paid for in cash). The Notes are in zero net supply.

During fifteen to twenty-five minutes, subjects can submit limit order which are posted in electronic books (one for each security); if a limit order crosses the best order at the other side of the market, it is automatically converted to a market order. Strict price and time priority are adhered to. No hidden limit orders are permitted. Order submission and trading are anonymous: only IDs (numbers between 100 and 200) are shown, but the identities behind IDs are never revealed.

After markets close, a state is drawn and dividends are determined depending on this state and the final holdings of the securities. Payments are made after subtracting a fixed amount referred to as "loan repayment," as discussed below. Subsequently, assets are taken away, and a new period starts (subjects are given a fresh supply of the state securities, as well as cash, etc.). Each experiment involved at least eight periods, in addition to a number of practice periods (the latter did not count towards total gain).

Subjects are to pay the experimenter for the securities and cash they are given at the beginning of each period. Effectively, this means that the experimenter gives securities and cash on loan, and hence, the payment is referred to as "loan repayment." The loan repayment creates leverage, causing a magnification of the risk involved in the holding of the state securities. It also means that subjects could lose money. A subject is barred from further trading if s/he has negative cumulative earnings for more than two periods in a row.[6]

All accounting is done in terms of a fictitious currency called francs, to be exchanged for dollars at the end of the experiment at a pre-announced exchange rate. In some experiments, subjects are also given an initial sign-up reward, which is fully exposed to risk (i.e., subjects run the risk of losing the entire sign-up reward during the experiment).

[5] We thank Nick Barberis for pointing this out to one of us.

[6] It can be shown that this *bankruptcy rule* induces risk-averse behavior even among risk-neutral and ambiguity-neutral subjects, except in the last period.

The relationship between states and payoffs, the *payoff matrix*, is the same for all experiments, namely:

State	X	Y	Z
Security X	100	0	0
Security Y	0	100	0
Security Z	0	0	100
Notes	100	100	100

The remaining data and parameters for the experiments are displayed in Table 1. There is variation in initial endowments across subjects, but always in such a way that the total ("social") endowment was such that security Y was in highest supply, and X is lowest supply. In other words, aggregate wealth was highest in state Y, and lowest in state X.

In one set of experiments, to be referred to as the *first set*, the probability of Z is announced; the relative probabilities of X and Y are not. That is, Z is the pure-risk state (corresponding to the extraction of a red ball in the classic Ellsberg 3-color experiment); X and Y are the ambiguous states (corresponding to the extraction of a blue respectively green ball). So, in the first set of experiments, the ambiguous securities are in extreme supply (aggregate wealth is highest and lowest in the ambiguous states). In the *second set* of experiments, the probability of X is announced, not those of Y and Z. This way, the ambiguous securities are both in higher supply than the pure-risk security.

Drawing of states is without replacement. In particular, we start the first period (that counts towards total payment) with an urn with 18 balls. As we draw balls (states), we do not replace them. This way, ambiguity is not reduced over the course of the experiment. When an ambiguous state is drawn, the subjects do not get any information about the relative composition of the urn as far as the ambiguous states are concerned, unless, of course, they start out with a well-defined prior, i.e., unless they are expected utility maximizers. As we shall document, some subjects fall in this category; they clearly updated from a uniform prior.

Note that there is substantial risk and uncertainty in our experiments. Subjects could and did lose money. About 3 subjects per experiment generated negative cumulative earnings during more than two periods in a row, and hence, were asked to leave the experiment. To induce competition, we organized our experiments at a relatively large scale, namely, about 30 subjects. This way, we could also afford to lose a few subjects to bankruptcy.

Except for experiment 021120, subjects were given oral instructions at the onset of an experiment, which they could follow on the experiment website. Experiment 021120 was not held at a central location with all subjects present, but decentralized, over the internet [like the pure-risk experiments discussed in (Bossaerts and Plott (2003), Bossaerts e.a. (2003))]. In that case, subjects had access to the instructions one day before the experiment and were invited to study the instructions. The instructions are the same throughout, except for the description of the drawing of the states, which is different across the two sets of experiments.

Experiments lasted about 3 hours (including practice sessions). Average per-subject payment in the experiments was approximately \$45; maximum pay was about \$125; minimum pay was \$0.

The websites for the experiments all have an URL with the same structure:

http://eeps3.caltech.edu/market-*,

where * should be replaced with the date of the experiment (see column 1 in Table 1). The interested reader can visit these websites, read the instructions (the typos were highlighted during the oral presentation), inspect the trading interface, and display the trading history (including pricing).[7]

3 Theoretical Predictions

Absent ambiguity aversion, 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 the level of aversion to ambiguity differs across agents, the effects are less obvious. This is because some agents may not anymore be marginal, effectively eliminating their influence on the relative pricing of ambiguous securities.

To understand how this works, we discuss three cases, gradually increasing complexity by adding agents with different perceptions and attitudes towards ambiguity. Highest complexity is obtained when we mix expected utility agents who update from uniform (diffuse) priors with ambiguity-neutral and ambiguity-averse agents.

As in the experiments, there are three states $s = r, g, b$.[8] 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 franc in state s where $s = j$ and zero otherwise.[9] 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.[10] 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 $1 - \pi$.

[7] Anonymous login requires the ID 1 and password a.

[8] In the experiments, the states are labeled X, Y and Z. To understand the connection with Ellsberg's paradox, we couch instead our theory in terms of "red," "blue," and "green." Later on, we map the results into the concrete X, Y, Z. The mapping will differ depending on which states are ambiguous. In the theory, states g and b are always ambiguous.

[9] In the experiments, state securities pay 100 francs. The necessary adjustment is straightforward.

[10] In the experiments, π changes over time.

3.1 Optimal Investments

3.1.1 Expected Utility Preferences

Suppose that subject $n \in \{1, ..., N\}$ 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. 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 condition for the optimal choice of state securities $w_s^{n,*}$, $s = r, g, b$, is standard:

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

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

For the sake of simple illustration later on, it is useful to introduce the following specialization of the above model.[11] Suppose that subject n's utility is quadratic with parameters $a_n > 0$ and $b_n > 0$:

$$u^n(w_s^n) = a_n w_s^n - \frac{b_n}{2}(w_s^n)^2.$$

As is well known, b_n measures subject n's risk aversion. Risk tolerance decreases as b_n increases. In the sequel, whenever considering this special case we make two assumptions. First, without loss of generality we assume $a_n = a$, all n. Second, we assume that no subject's satiation point is ever reached.

In this quadratic case the resulting optimal demands exhibit a property that is known as *portfolio separation*: all subjects' demands can be decomposed into the demand for a riskfree security (say, one that pays 1 in each state) and demand for a common risky security (one with payoff p_s/π_s in state s):

$$\begin{bmatrix} w_r^{n,*} \\ w_g^{n,*} \\ w_b^{n,*} \end{bmatrix} = \frac{a}{b_n}\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \frac{\lambda^{n,*}}{b_n}\begin{bmatrix} \frac{p_r}{\pi_r} \\ \frac{p_g}{\pi_g} \\ \frac{p_b}{\pi_b} \end{bmatrix}.$$

[11] On the other hand, it should be observed that, since the risk in our experiments is fairly small for the subjects, quadratic preferences do offer a good local approximation of subjects' actual utility profiles. Under pure risk, the appropriateness of the quadratic approximation is confirmed in the experimental data. See (Bossaerts e.a. (2003)).

3.1.2 α-Maxmin Preferences With $\alpha \geq 0.5$

Consider next a subject n who perceives ambiguity in the portfolio choice problem, and whose attitudes towards risk and ambiguity that can be summarized in terms of α-maxmin quadratic preferences (see (Ghirardato e.a. (2004))), with the following set of priors:[12]

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

As for an expected utility subject, the number π gets updated according to Bayes' law as information is received on the urn composition.[13] Subject n's total utility in this case 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 the (state-independent) utility function u^n is again assumed to be twice differentiable, increasing and strictly concave.

Intuitively, the coefficient α_n measures subject n's ambiguity aversion. 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.[14]

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

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

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

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

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

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

[12] Δ^2 denotes the two-dimensional unit simplex in $\mathbb{R}^3$.

[13] If, 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.

[14] This "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.

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.

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:

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

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.

3.2 Equilibrium

3.2.1 Equilibrium without Ambiguity Attitudes

Consider first the case in which all agents have expected utility preferences as described above, and use the prior $[\pi, \rho, 1-\pi, \rho]$ (in our experimental setup, this would naturally be the uniform prior $[1/3, 1/3, 1/3]$).

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 this case it is well known (see, e.g., (Bossaerts e.a. (2003))) 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.

To see this very clearly, consider the special case of subjects with quadratic utility functions discussed earlier. Let W_s^0 denote the per-capita supply of state security s, i.e.,

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

The equilibrium state price probability ratios are readily obtained:

$$\frac{p_s^*}{\pi_s} = \frac{\frac{1}{N} a \sum_{n=1}^{N} \frac{1}{b_n}}{\frac{1}{N}\sum_{n=1}^{N} \frac{\lambda^{n,*}}{b_n}} - \frac{1}{\frac{1}{N}\sum_{n=1}^{N} \frac{\lambda^{n,*}}{b_n}} W_s^0.$$

This can be simplified further, because individual subjects' first-order conditions imply that[15]

$$\frac{\lambda^{n,*}}{b_n} = \frac{a}{b_n} - \sum_s \pi_s w_s^{n,*}. \tag{6}$$

Consequently,

$$\frac{1}{N}\sum_{n=1}^{N} \frac{\lambda^{n,*}}{b_n} = \frac{a}{B} - E[W^0],$$

where B denotes the *harmonic mean risk aversion*

$$B = \left(\frac{1}{N}\sum_{n=1}^{N} \frac{1}{b_n}\right)^{-1},$$

and $E[W^0]$ is the expected per-capita payout,

$$E[W^0] = \sum_s \pi_s \frac{1}{N}\sum_{n=1}^{N} W_s^0.$$

Altogether:

$$\frac{p_s^*}{\pi_s} = \frac{a}{a - BE[W^0]} - \frac{B}{a - BE[W^0]} W_s^0. \tag{7}$$

That is, in equilibrium there exists a representative agent whose utility can be chosen to be quadratic as well.[16] Moreover, the larger W_s^0, the smaller the ratio p_s^*/π_s.

3.2.2 Equilibrium with Ambiguity Attitudes

Assume next that subjects are sensitive to the ambiguity in the problem, so that subjects' preferences and portfolio choices are as described in Section 3.1.2 (with common set of priors C). If $\alpha_n = 1$ for all n, then it is immediate

[15] To see this, sum across the first-order conditions and use the restrictions that $\sum_s \pi_s = 1$ as well as $\sum p_s = 1$.

[16] We choose this agent's risk aversion coefficient to be B; we set the coefficient of the linear term of the quadratic utility function, A, and the Lagrange multiplier, Λ^*, to values that satisfy the first-order conditions.

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).[17]

If, on the other hand, $\alpha_n = \alpha \in [0.5, 1)$ for all n, then it is again 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 in the first period 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), while in latter periods is given by the equilibrium that obtains with beliefs $[\pi', (1-\alpha)(1-\pi'), \alpha(1-\pi')]$ with π' obtained from π by application of Bayes' rule. 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.[18] 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.

On the other hand, as we recalled earlier, 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, it seems compelling to consider next a "mixed" situation where subjects may display heterogeneity in such characteristics.

So, suppose that the subjects are of two types. They can either be ambiguity-insensitive expected utility maximizers, with prior beliefs given by $[\pi, \rho, 1-\pi-\rho]$, or be extremely ambiguity averse and have 1-maxmin preferences with the set of priors C. (It is simple to see how the analysis to follow is adapted to the case in which expected utility is substituted with 1/2-maxmin.)

Intuitively, Figure 1 demonstrated that, for a given level of p_g/p_b, α-maxmin subjects will be exposed to conditional risk in the ambiguous states only if their α falls below a critical value. One would therefore conjecture that only subjects with expected utility preferences 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.

To determine the equilibrium pricing restrictions explicitly, suppose that all subjects have quadratic utilities and order them so that subjects $n = 1, ..., N^*$ ($N^* < N$) have expected utility preferences and subjects $n > N^*$ have

[17] The nonnegativity of prices is the reason behind the apparent discrepancy with the result of (Epstein and Wang (1994)) that an equilibrium exist (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.

[18] This 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.

α-maxmin preferences with $\alpha_n = 1$. The equilibrium condition for the market of the pure-risk state security r is standard:

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

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

$$\begin{array}{rcl} \sum_{n=1}^{N^*} w_g^{n,*}(p^*) & = & \sum_{n=1}^{N} w_g^{0,n} - \sum_{n=N^*+1}^{N} w_g^{n,*}(p^*), \\ \sum_{n=1}^{N^*} w_b^{n,*}(p^*) & = & \sum_{n=1}^{N} w_b^{0,n} - \sum_{n=N^*+1}^{N} w_g^{n,*}(p^*), \end{array} \tag{8}$$

where we used the fact that $w_g^{n,*} = w_b^{n,*}$ for $n > N^*$. [19]

The *equilibrium state price probability ratio* for the risky state security is as before:

$$\frac{p_r^*}{\pi} = \frac{\frac{1}{N}\sum_{n=1}^{N}\frac{a}{b_n}}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} - \frac{1}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} W_r^0.$$

The equilibrium state price probability ratios of the ambiguous state securities are more complicated. It is most useful to write them as follows.

$$\frac{p_g^*}{\rho} = \left(\frac{\frac{1}{N}\sum_{n=1}^{N}\frac{a}{b_n}}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} + \frac{\frac{1}{N}\sum_{n=N^*+1}^{N}\frac{\lambda^{n,*}}{b_n}}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}}\frac{p_g - p_b}{1-\pi}\right) - \frac{1}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} W_g^0,$$

$$\frac{p_b^*}{1-\pi-\rho} = \left(\frac{\frac{1}{N}\sum_{n=1}^{N}\frac{a}{b_n}}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} - \frac{\frac{1}{N}\sum_{n=N^*+1}^{N}\frac{\lambda^{n,*}}{b_n}}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}}\frac{p_g - p_b}{1-\pi}\right) - \frac{1}{\frac{1}{N}\sum_{n=1}^{N}\frac{\lambda^{n,*}}{b_n}} W_b^0.$$

One can simplify the equilibrium conditions for the prices the way we did before, because (6) holds here as well. So,

$$\frac{p_r^*}{\pi} = \frac{a}{a - BE[W^0]} - \frac{B}{a - BE[W^0]} W_r^0;$$

$$\frac{p_g^*}{\rho} = \left(\frac{a}{a - BE[W^0]} + B\frac{\frac{1}{N}\sum_{n=N^*+1}^{N}\frac{\lambda^{n,*}}{b_n}}{a - BE[W^0]}\frac{p_g - p_b}{1-\pi}\right) - \frac{B}{a - BE[W^0]} W_g^0;$$

$$\frac{p_b^*}{1-\pi-\rho} = \left(\frac{a}{a - BE[W^0]} - B\frac{\frac{1}{N}\sum_{n=N^*+1}^{N}\frac{\lambda^{n,*}}{b_n}}{a - BE[W^0]}\frac{p_g - p_b}{1-\pi}\right) - \frac{B}{a - BE[W^0]} W_b^0.$$

This way, the same coefficient appears in front of the per-capita supply in each of the three equilibrium conditions. As a consequence, we can determine the effect of changes in the relative supplies of the three state securities on the pattern of state price probability ratios. Indeed, as a function of supplies, the three equilibrium conditions represent

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

straight lines with different intercepts but common slope. The slope depends only on the average per-capita supply across states. At the same time, the intercepts depend in complicated ways on the distribution of equilibrium Lagrange multipliers as well – but the distances between the intercepts for the ambiguous states and that for the pure-risk state are the same. So, the equilibrium conditions for the ambiguous securities can be interpreted as equidistant, positive or negative vertical translations of the equilibrium line of the pure-risk security.

In the absence of heterogeneity (in attitudes towards ambiguity), equilibrium state price probability ratios for ambiguous securities would be on the (single) equilibrium line for the risky security. Because of heterogeneity, they are now displaced an equal distant vertically below and above this line. As a consequence, the *range and variability of the state price probability ratios* is increased, potentially generating rankings that are incompatible with expected utility maximization.

To understand this effect in more detail, *first*, assume that $W_g^0 < W_r^0 < W_b^0$. That is, the supplies of ambiguous securities are extreme: one is in short supply, the other is abundant. The supply of the risky security is moderate. Of course, for this supply pattern to be held in equilibrium, it must be that $p_g > p_b$. Inspection of the three equilibrium conditions then suggests that the state price probability ratios will be on three parallel lines, with the line representing the equilibrium condition for the risky state (state r) in the middle, and the other two lines equidistant (vertically) above (state g) and below (state b) it. See Figure 2.

Figure 2 demonstrates that the *ranking of the state price probability ratios will always be in accordance to theoretical predictions if there were only risk aversion*: the ranking is inversely related to per capita endowment of the state securities. But ignoring the presence of ambiguity averse subjects, however, will lead to an overstatement of risk aversion, as we now explain.

For there to be a representative agent with state-independent quadratic utility, one must be able to connect the circles (representing the equilibrium state price probability ratios) with a straight line. This line represents the marginal utility of the representative agent; it is indeed to be linear if the representative agent is to have state-independent quadratic utility. Figure 2 suggests that in general, it is impossible to connect the circles with a single straight line; i.e., in general, there will *not be a representative agent with state-independent quadratic utility.* However, the error will usually be small (compare to the cases below).

If we were to fit a straight line through the circles, it is obvious that it will have a slope that is *bigger* (in absolute value) than that of the three lines representing the equilibrium pricing conditions. In other words, the *risk aversion implied by the state price probability ratios is over-stated*: it reflects in part the presence of ambiguity averse agents.

Consider, *second*, the case where $W_r^0 < W_g^0 < W_b^0$. In other words, the ambiguous state securities are in bigger supply. *Perverse ranking* of state price probability ratios now becomes possible (perverse relative to predictions ignoring ambiguity aversion). Figure 3 provides an example. This figure also illustrates that one is much more likely to reject state-independent quadratic utility for the representative agent: a linear fit through the circles will generate a more sizeable error than in the first case.

Thirdly, consider the case where $W_g^0 < W_b^0 < W_r^0$, which means that the pure risk state security is in higher supply than the ambiguous state securities. Figure 4 plots the equilibrium pricing lines. Again, *perverse ranking* becomes possible (perverse relative to a model that ignores the presence of ambiguity averse subjects). Utility of the representative agent has to be made state-dependent, as in the second case.

Notice that if there is *correlation between risk aversion and ambiguity aversion* in the sense that subjects $n = N^* + 1, ..., N$ have a higher b_n, then this will tend to *decrease* the distance between the pricing lines. As a result, state price probability ratios will look more as if the representative agent had state-independent quadratic utility.

3.2.3 Heterogeneous α's and Expected Utility

We close our discussion of the equilibrium predictions by briefly considering a further differentiation of the subjects, into three categories. Precisely, we suppose that agents can either be expected utility maximizers or have α-maxmin preferences with $\alpha \in \{0.5, 1\}$. As we shall see presently, this extension has the advantage of explaining a particular feature of our data.

This situation could be rather complex, for instance, when $W_g^0 > W_b^0$ and π_g (to be interpreted as the belief of expected utility subjects) $> (1/2)(1-\pi)$, or $W_g^0 < W_b^0$ and $\pi_g < (1/2)(1-\pi)$. What we have here is a mixture of expected utility maximizers who think state g has more than probability $(1-\pi)/2$, whereas ambiguity "neutral" subjects ($\alpha = 0.5$) effectively assign probability $(1-\pi)/2$ to state g. For simplicity, let us assume that there are very few ambiguity neutral subjects, so that prices of ambiguous securities are driven by the expected utility maximizers. The issue is: what positions would the ambiguity neutral subjects take?

The positions of the (price-setting) expected utility maximizers should be clear: they accommodate the inequality in supplies of ambiguous securities and buy more of security g than of security b. To induce them to do so, the state price probability ratio for state g should be less than that of state b, where the state price probability ratios are determined based on their beliefs. Because $\pi_g > (1-\pi)/2$, it is conceivable that $p_g > p_b$. But this induces the ambiguity neutral subjects to invest *more in state b than in state g* (see (2)), exactly the opposite of the expected utility maximizers! Consequently, it is possible to see three different kinds of portfolios:

1. Expected utility maximizers invest more in state g than in state b;

2. Ambiguity averse subjects invest an equal amount in states g and b;

3. Ambiguity neutral subjects invest less in state g than in state 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 exposure to risk and holdings of ambiguous securities 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.

4.1 State Price Probability Ratios

In experiments 030123, 030203 and 030210, state securities X and Y were ambiguous and they were the securities in lowest and highest supply, respectively. Figure 5 displays the empirical distribution function of the three 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. These are the state price probability ratios that we used to analyze pricing when the population is either homogeneous expected utility (with a common, uniform prior – Section 2.1) or heterogeneous expected utility (with a uniform prior) and maxmin (Section 2.5).

We display empirical distribution functions of the state price probability ratios because they 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 subjects did exhibit aversion against ambiguity.[20] Because of the sizeable bid-ask spreads, volatility remained high throughout the experiment.

Figure 5 reveals a common pattern to the three experiments where X and Y were ambiguous states: the probability that X exceeded any level is less than that for Z, which itself is less than that of Y. One can express this in terms of quantiles: the median (or any other quantile) of the state price probability ratio of X is highest; that of Y is

[20] Under 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.

lowest. This conforms to the prediction when subjects have expected utility preferences, and hence, are insensitive to ambiguity. It does correspond as well, however, to the case where there is heterogeneity, but were the expected utility subjects set the relative prices of the two ambiguous states (X and Y).

The situation could, however, be as depicted in Figure 2, where the ordering of state price probability ratios is inversely related to aggregate wealth as if only expected utility subjects are present, but there nevertheless are ambiguity averse subjects. The only way to ascertain the latter is to evaluate the level of risk aversion reflected in pricing, but we refrain from doing so here, because it would ask for inter-experiment comparisons of risk premia, a rather tricky task. Instead, inspection of subjects' final holdings (see below) will confirm that some subjects were indeed ambiguity averse, and hence, that Figure 3 is the right picture with which to interpret the pricing results.

We deflated state prices with probabilities based on a uniform prior (over the ambiguous states) and updating according to Bayes' law. One wonders to what extent Figure 5 is sensitive to alternative assumptions about these probabilities. One important possibility is that the marginal investor for the ambiguous securities is α-maxmin with $\alpha = 0.5$ (see Section 2.3). This corresponds to expected utility maximization from a uniform prior without updating. Unfortunately, because of the actual sequence of draws, the empirical distribution functions of the state price probability ratios generated under this alternative are hardly any different from the ones depicted in Figure 5. This happened not to be the case, however, in some of the experiments where Y and Z (securities in highest supply) were ambiguous. Let us discuss those now.

Figure 6 displays the empirical distribution functions of the state price probability ratios in the four experiments where Y and Z (the two securities in highest supply) were ambiguous. To deflate state prices, probabilities were used that were updated from a uniform prior. In experiments 021120 and 030219, the ordering of state prices probability ratios is as if no ambiguity aversion is present: the median (or any quantile) of X is highest; that of Y is smallest. Yet the ordering of state price probability ratios of X and Z in 020529 and 020619 is anomalous (relative to expected utility): in either case, there is no clear first-order stochastic dominance of the state price probability ratio of Z over that of X.

The latter is precisely the situation depicted in Figure 3, where the two state securities in highest supply are also the ambiguous securities, but where there is heterogeneity in attitudes towards ambiguity, so that more ambiguity averse subjects do not determine the relative pricing of the ambiguous securities. Consequently, there appears to be heterogeneity in attitudes towards ambiguity. It shows up in a distinct pattern in state price probability ratios when the two state securities in highest supply are the ambiguous securities: the state price probability ratios of the poorest state and that of the middle state can hardly be ranked relative to each other. Again, we are going to confirm the presence of heterogeneity in ambiguity aversion by studying the final holdings of individual subjects.

Before we do so, let us discuss the effect of a different deflation of state prices, namely, using a fixed, uniform probability for the ambiguous states. We thereby assume that the marginal investor in the markets for ambiguous securities is α-maxmin with $\alpha = 0.5$. Figure 7 displays the results. A distinct anomaly emerges in three of the four experiments: the state price probability ratios of the two ambiguous states, Y and Z, do not exhibit a clear stochastic dominance relationship anymore. For all practical purposes, the state price probability ratios of Y and Z are equal, as if the ambiguity-neutral subjects demand less of a risk premium for those securites. The intuition behind the paradoxical nature of this finding is simple: a few (ambiguity neutral) subjects are in fact asked to hold Y and Z in far more unbalanced ways than if everybody had been ambiguity neutral, because the ambiguity averse subjects choose to hold them in equal proportions. If anything, the ambiguity neutral subjects can be expected to demand a higher risk premium to absorb the unequal supplies of Y and Z that others refuse to accommodate.

Still, as we will show next, there is some evidence that the ambiguity neutral subjects are in fact less risk averse, tempering some of the price adjustment that is needed for them to willingly accommodate the unequal supplies of ambiguous securities.

4.2 Final Holdings

Prices have given us some indication that there is heterogeneity in attitudes towards ambiguity. End-of-period holdings confirm this. Furthermore, these holdings reveal correlation between ambiguity and risk aversion.

Figure 8 plots (i) the proportion of end-of-period wealth exposed to ambiguity allocated to the ambiguous security in highest supply (Y), against (ii) the volatility of end-of-period wealth. The latter is measured as the standard deviation of a subject's wealth across states, assuming a uniform (but updated) prior over the ambiguous states. It is a crude indication of a subject's level of risk tolerance.[21]

The figure clearly indicates that the majority of subjects have little or no exposure to ambiguity, and that only a minority of them hold ambiguous securities in very unbalanced numbers (sometimes managing to short some of security X in order to invest in Y). The figure also reveals a remarkable correlation between ambiguity tolerance (willingness to hold Y as a proportion of X and Y different from 0.5) and risk tolerance (as measured by ex ante volatility of end-of-period wealth).

Heterogeneity in ambiguity aversion and positive correlation between ambiguity and risk aversion also showed up in the other four experiments. See Figure 9. The results are a bit more noisy now, not surprisingly, because the difference in supplies of the two ambiguous securities (Y and Z) is not anymore as large as it was in the previous three experiments (the results of which where depicted in Figure 8). Ambiguity tolerant subjects are not required to do as much accommodating.

Note, however, that there are a fair number of subjects in 020529, 020619 and 021120 who invested *less in security Y than in Z*. We discussed this possibility in the theoretical section. Basically, it would emerge if there are both

[21] Volatility is measured in francs.

ambiguity neutral subjects and expected utility maximizers, and the price of the security in highest supply (Y) has increased above that of the other security (Z). Figure 7 shows that there was a high likelihood for Y to be more expensive than Z in these experiments.[22]

The latter illustrates the complex nature of the effects of heterogeneity: the beliefs of the expected utility maximizers drive the prices of the ambiguous securities in a direction that entices ambiguity neutral subjects to buy more of the ambiguous security that is in lowest supply.

5 A Final Remark

The observed violations of the basic relationship between the rankings of state price probability ratios and aggregate wealth and even the patterns between holdings of ambiguous securities and risk exposure can be explained not only by heterogeneity in attitudes towards ambiguity, but also by heterogeneity in beliefs (about the relative likelihood of the ambiguous states). There is no *a priori* reason, however, why the distribution of beliefs changes in such a way as to mimic the effects of changes in relative supplies of the securities predicted under heterogeneous ambiguity aversion. In this sense, the violations of standard equilibrium rankings observed in our experiments can be explained most parsimoniously in terms of heterogeneity in ambiguity aversion.

Further experiments are needed to differentiate more sharply between heterogeneity in beliefs and heterogeneity in ambiguity aversion as explanations of violations of standard equilibrium ranking predictions of state price probabilities and patterns of final holdings across subjects. These may have to involve changes in aggregate supplies of state securities within an experiment. We are presently designing such experiments, to be reported on in a sequel paper.

References

Bossaerts, P. and C. Plott (2003): “Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets.” Caltech working paper.

Bossaerts, P., C. Plott and W. Zame (2003): “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.

[22]The figure shows state pice probability ratios, not prices. But these are obtained by deflating the prices of Y and of Z by the same probability, namely, $(1-\pi)/2$, so that the relationship between prices of Y and Z can readily be derived from the relationship between the state price probability ratios.

Dana, R.A. (2004): "Ambiguity, Uncertainty Aversion and Equilibrium Welfare." *Economic Theory* 23, 569-587.

Ellsberg, D. (1961): "Risk, Ambiguity and the Savage Axioms." *Econometrica* 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.

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.

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.

Maenhout, P. (2000): "Robust Portfolio Rules and Asset Pricing." INSEAD working paper.

Uppal, R. and T. Wang (2003): "Model Misspecification and Under Diversification." *Journal of Finance*, forthcoming.

Table 1: Experiment Parameters

Date	Urn (18 Total)			Subject Category	Signup Reward	Endowments[a]			Loan	Exchange Rate
	X	Y	Z	(Number)	(franc)	X	Y	Z	(franc)	cents/franc
030123	?	?	9	13	100	4	11	4	500	2
				13	100	1	4	6	375	2
030203	?	?	9	15	500	4	11	4	500	2
				14	500	1	4	6	375	2
030210	?	?	6	15	500	4	11	4	500	2
				14	500	1	4	6	375	2
020529	6	?	?	13	0	4	11	2	220	2/3
				13	0	1	4	8	300	2/3
020619	6	?	?	12	0	4	11	2	220	2/3
				12	0	1	4	8	300	2/3
021120	6	?	?	15	0	4	11	2	220	2/3
				14	0	1	4	8	300	2/3
030219	6	?	?	15	300	4	11	2	500	2
				14	300	1	4	8	375	2

[a] Number of each of the state securities, in addition to 200 francs cash. No net holdings of Notes.

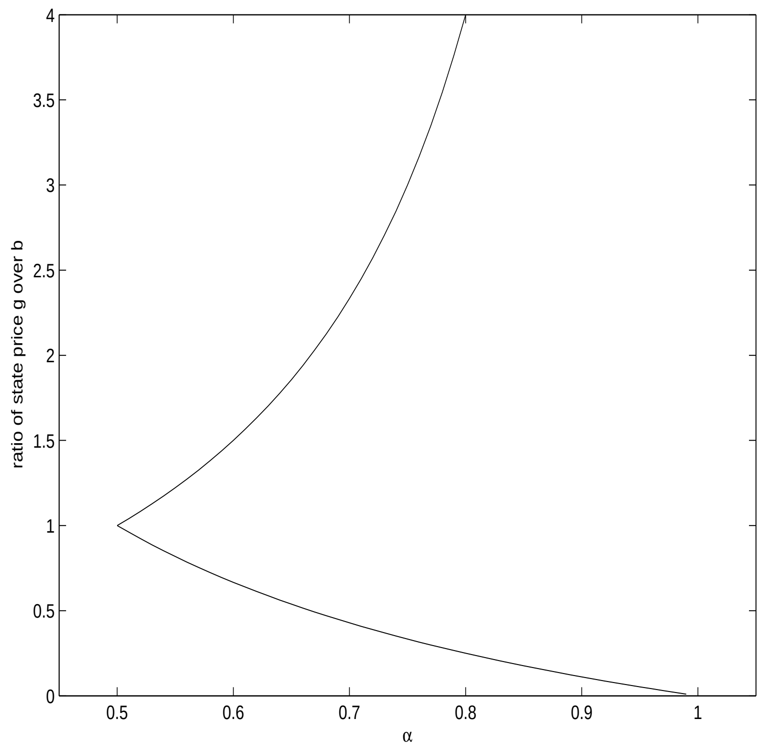

Figure 1: Ratio of the prices of security g over b above which an α-maxmin agent sells security g (top line) or below which (s)he buys (bottom line). α is on the X axis; $\alpha = 0.5$ corresponds to ambiguity neutrality; $\alpha = 1$ corresponds to extreme ambiguity aversion.

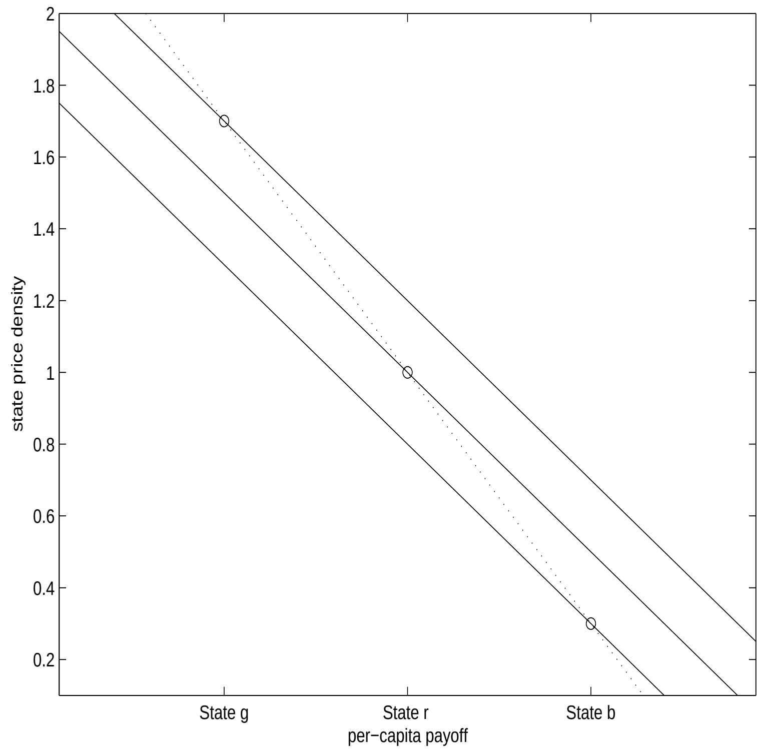

Figure 2: Equilibrium state price density (i.e., state price probability ratios – see circles) when the ambiguous state securities are in lowest (state security g) and highest (state security b) supply. (State r is the risky state.) In the situation that is depicted here, the state price density is on a single straight line (dotted), as if there were a representative agent with state-independent quadratic utility, albeit with risk aversion (slope) higher than that of the actual agents (reflected in the slopes of the parallel, solid lines).

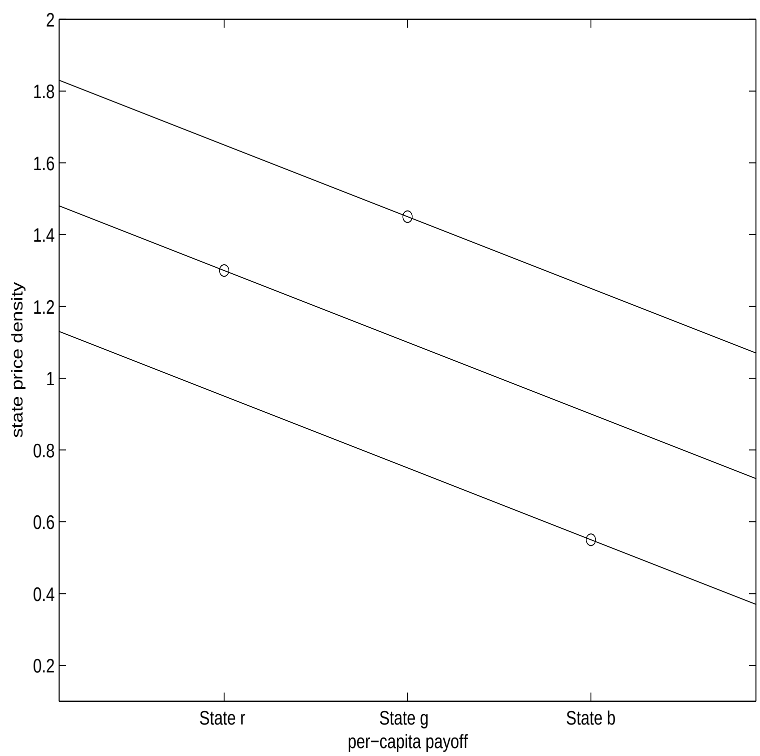

Figure 3: Equilibrium state price density (state price probability ratios – see circles) when the ambiguous state securities (state securities g and b) are in highest supply. (State security r, in lowest supply, pays in the pure-risk state.) The state price probability ratios cannot anymore be connected by a single straight line, as if there were a representative agent with state-dependent quadratic utility. Note that the state price probability ratio of the pure-risk state (state r) can now be below that of the ambiguous state security in lowest supply (state g). This situation will never occur under homogeneous expected utility preferences.

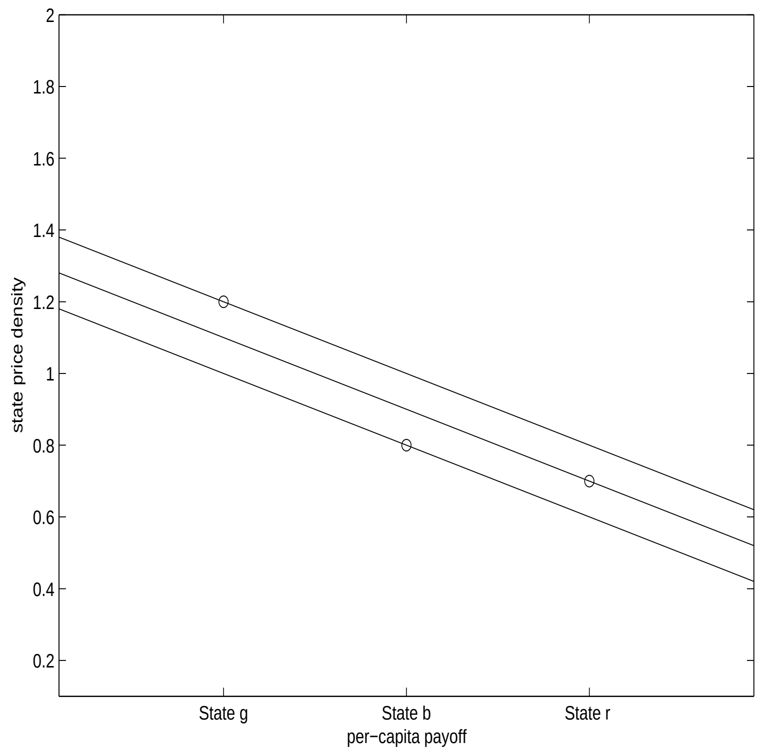

Figure 4: Equilibrium state price density (circles) when the ambiguous state securities (state securities g and b) are in lowest supply. (State security r, in highest supply, pays in the pure-risk state.) The state price probability ratios cannot be connected by a single straight line, as if there were a representative agent with state-dependent quadratic utility. Note that the state price probability ratio of the pure-risk state (state r) could now be above that of the ambiguous state security in highest supply (state b) (this situation is not depicted; it cannot occur under homogeneous expected utility preferences).

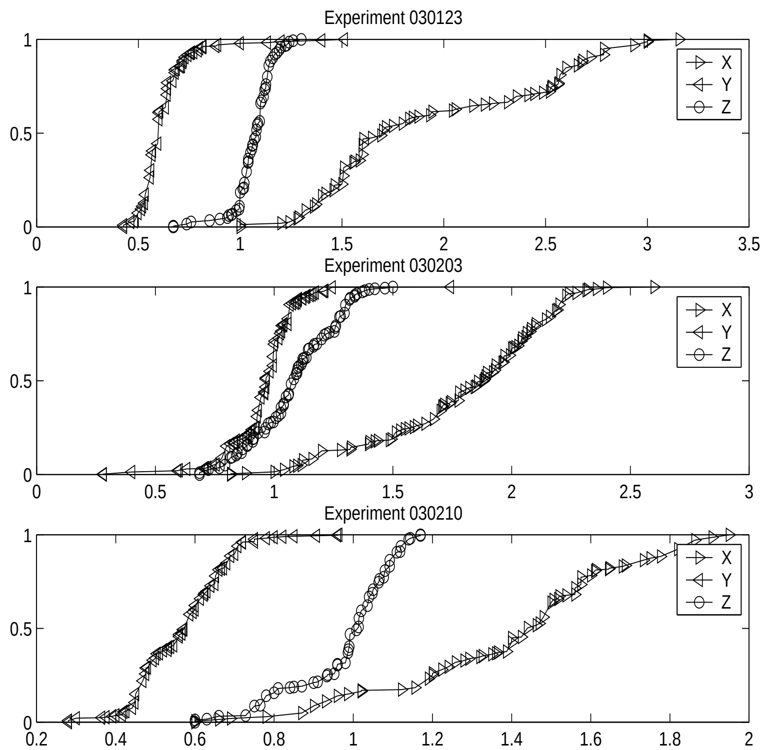

Figure 5: Empirical distribution functions of state price probability ratios in three experiments where the ambiguous securities, X and Y, are in lowest and highest supply, respectively. State price probability ratios are obtained by deflating state prices with probabilities based on (i) a diffuse prior over ambiguous states, (ii) updated using Bayes' law and the draws of the states in prior periods. As long as there are expected utility maximizing subjects, the state price probability ratios will tend to be highest for X and lowest for Y; the presence of ambiguity averse subjects is only discernible in an exacerbated risk premium (reflected in unusually large differences between the state price probability ratios of X, Y and Z).

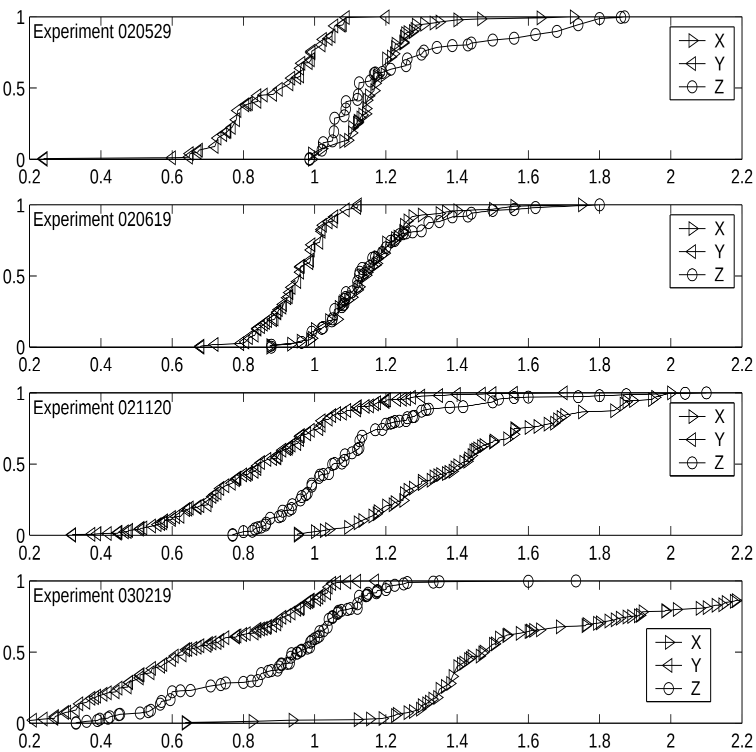

Figure 6: Empirical distribution functions of state price probability ratios in four experiments where the ambiguous securities, Y and Z, are in highest supply. (The pure risk state security, X, is in lowest supply.) State price probability ratios are obtained by deflating state prices with probabilities based on (i) a diffuse prior over ambiguous states, (ii) updated using Bayes' law and the draws of the states in prior periods. As long as there are expected utility maximizing subjects, the state price probability ratios will tend to be highest for X and lowest for Y; the presence of ambiguity averse subjects is discernible in potential reversal in ordering between the state price probability ratios of the ambiguous state security in lowest supply (Z) and that of the pure risk state (X). See experiments 020529, 020619 and 021120.

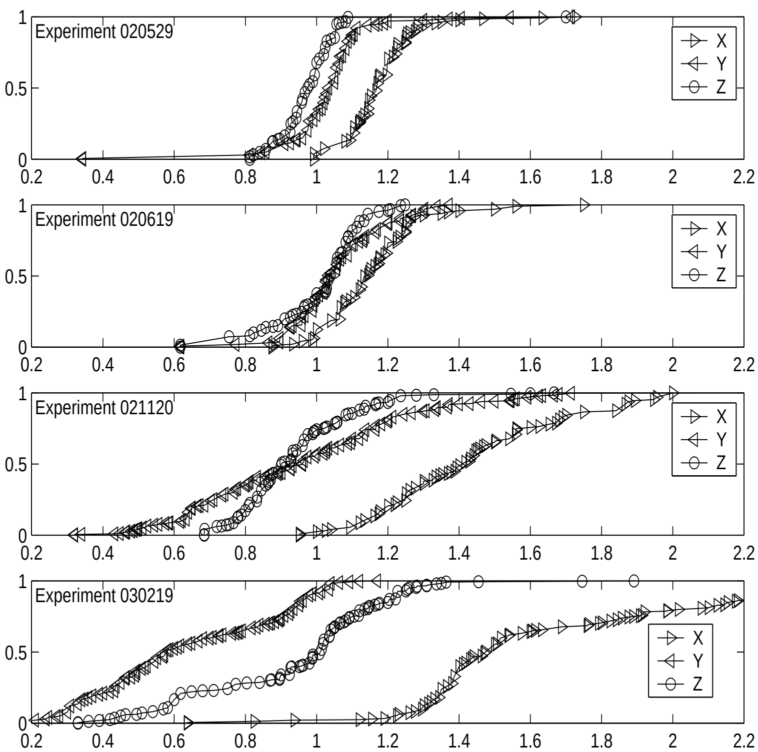

Figure 7: Empirical distribution functions of state price probability ratios in four experiments where the ambiguous securities, Y and Z, are in highest supply. (The pure risk state security, X, is in lowest supply.) State price probability ratios are obtained by deflating state prices with probabilities whereby only the probability of the risky state (π) is updated; those of the ambiguous states remain $(1-\pi)/2$. This deflation generates a tendency of anomalous orderings of state price probability ratios: state Y is more expensive than state Z, which would entice ambiguity neutral subjects to hold more of Z than of Y, in contrast to what is needed for equilibrium (Y is in higher supply than Z).

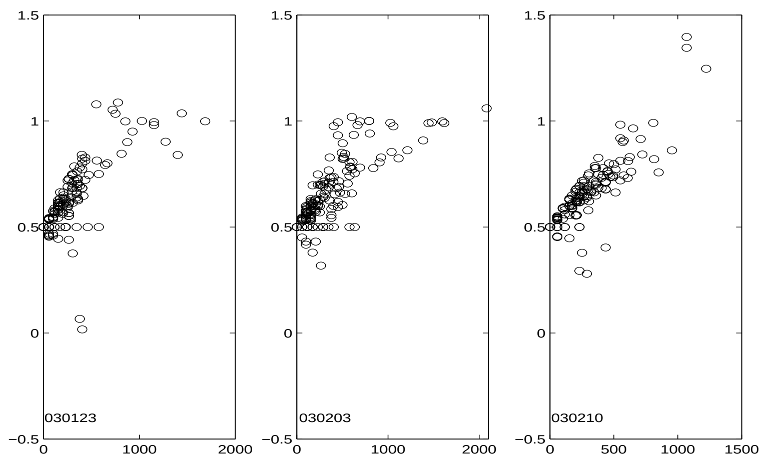

Figure 8: Plot of holdings of the ambiguous security in highest supply (Y) as a proportion of holdings of all ambiguous securities (X and Y) against volatility (in francs) of end-of-period individual wealth (horizontal axis) in three experiments where states X and Y were ambiguous. Ambiguity averse subjects tend not to be exposed to risk in the ambiguous states, which means that they will hold the ambiguous securities in equal proportion (0.5). Ambiguity neutral and expected utility maximizers will be enticed, through pricing, to hold more of Y than of X. The horizontal axis measures subjects' tolerance for risk. Note the positive correlation between risk tolerance (horizontal axis) and ambiguity tolerance (vertical axis).

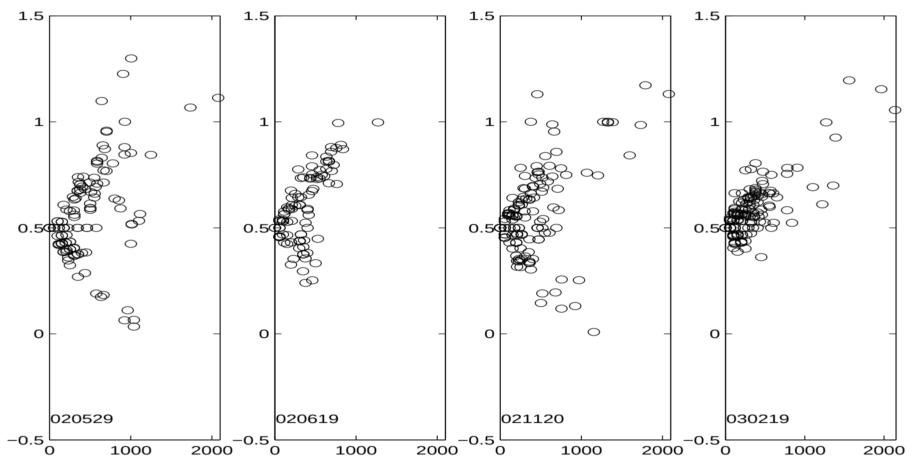

Figure 9: Plot of holdings of the ambiguous security in highest supply (Y) as a proportion of holdings of all ambiguous securities (Y and Z) against volatility (in francs) of end-of-period individual wealth (horizontal axis) in three experiments where states Y and Z were ambiguous. Ambiguity averse subjects tend not to be exposed to risk in the ambiguous states, which means that they will hold the ambiguous securities in equal proportion (0.5). Ambiguity neutral and expected utility maximizers will be enticed, through pricing, to hold more of Y than of Z. Sometimes, however, the pricing necessary to induce expected utility maximizers to do so leads the ambiguity neutral subjects to invest in the opposite direction. Whence the cases whereby subjects hold *less of Y than of Z* (see the circles substantially below 0.5 in experiments 020529, 020619 and 021120). The horizontal axis measures risk tolerance.