

Ambiguity and Asset Prices: An Experimental Perspective

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

This version: January 2006

[‡]Bossaerts: Caltech and CEPR; Guarnaschelli: McKinsey; Ghirardato: Università di Torino & Collegio Carlo Alberto; Zame: UCLA.

Ambiguity and Asset Prices: An Experimental Perspective

Abstract

Most of the economics and finance literature assumes that individual agents obey the Savage axioms; that is, they maximize expected utility according to subjective priors. However, Knight, Ellsberg and others argue that individual agents distinguish between risk (known probabilities) and uncertainty, or ambiguity (unknown probabilities), and that individual agents may display aversion to ambiguity, just as they display aversion to risk.

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: 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 densities. Experiments confirm the theoretical predictions.

The experiments also suggest a positive correlation between risk aversion and ambiguity aversion, which in turn suggests an explanation of the “value effect.”

Ambiguity and Asset Prices: An Experimental Perspective[†]

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 marginal rates of substitution, 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 leads to a wider range of state price densities, to reversals of ranking of state price densities, and to asset pricing that is *inconsistent* with the existence of a representative agent who maximizes expected utility.

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 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

[†]We are grateful for comments to seminar audiences at CIRANO and UC Irvine. 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), and the UCLA Academic Senate Committee on Research (Zame). 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).

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.

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,

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.

state price densities (ratios of state prices to probabilities) are necessarily ordered in the opposite direction as supplies of state securities. (That is, the state price density 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 (including expected utility maximizers). This effect may distort the state price densities.⁵ This distortion is likely to be largest when ambiguous securities are in greatest supply or least supply. If the distortion is small, state price densities will still be ordered in the direction opposite to supply; if the distortion is greater, state price densities may not be ordered in the direction opposite to supply. In both cases, this is exactly what we see in our experimental markets. And 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. Again, the data from our experimental markets bear out this prediction.

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

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

⁵In the presence of ambiguity, we define state price densities using a uniform prior over ambiguous states.

⁶We thank Nick Barberis for pointing this out.

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.

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, Knightean 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)]. Our results suggest a problem with this literature, since heterogeneity in ambiguity attitudes leads (even in relatively small groups) leads to pricing that are not rationalizable with a representative agent.

Following this Introduction, Section 2 describes our experimental design and Section 3 presents a theoretical analysis, generating asset pricing and portfolio holding predictions. Section 4 analyzes the data in view of these predictions. Section 5 concludes with a final remark about asset pricing with expected utility and 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.⁷

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).

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. In the *third set*, ran more than one year after the other two sets (to avoid contamination), we replicated on experiment from the first and one from the second set *perfectly*, but *without ambiguity*. That is, we used the same per-capita total supplies of

⁷It can be shown that this *bankruptcy rule* induces risk-averse behavior even among risk-neutral and ambiguity-neutral subjects, except in the last period.

all the securities,⁸ and, most importantly *exactly the same draws*. We informed the subjects about the composition of the urn; it was the composition that was used to draw the outcomes in the original experiments.

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 (when present) 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 ?]. 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).⁹

3 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;

⁸The number of subjects in the experiments was also approximately the same. In one experiment, we decreased the signup reward (relative to the original experiment); in the other experiment, we increased the exchange rate and the signup reward.

⁹Anonymous login requires the ID 1 and password a.

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.

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 $s = r, g, b$.¹⁰ 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.¹¹ 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 .¹² 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 $1 - \pi$.

3.1 Optimal Investments

3.1.1 Expected Utility Preferences

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 condition for the optimal choice of state securities $w_s^{n,*}$, $s = r, g, b$, is standard:

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

¹⁰In 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.

¹¹In the experiments, state securities pay 100 francs. The necessary adjustment is straightforward.

¹²In the experiments, π changes over time.

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 a specialization of the above model in which all agents' utilities have a quadratic form.¹³ 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 $a_n = a$, for 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):

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

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

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:¹⁴

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

As for an expected utility subject, the number π gets updated according to Bayes' law as information is received on the urn composition.¹⁵ 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

¹³ 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, so that this specialization is less restrictive than it may appear at first glance. Under pure risk, the appropriateness of the quadratic approximation is confirmed in the experimental data. See Bossaerts e.a. (2005).

¹⁴ Δ^2 denotes the two-dimensional unit simplex in \mathbb{R}^3 .

¹⁵ 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 $5/16$.

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.¹⁶

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

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

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

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

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

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

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.

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:

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

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.

¹⁶This ‘‘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.

3.2 Equilibrium

3.2.1 Equilibrium with Heterogeneous Risk Attitudes and without Ambiguity Attitudes

Consider first the case in which all agents have expected utility preferences as described above (with possibly different utility functions), 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. (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.

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 that¹⁷

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

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*

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

¹⁷To see this, sum across the first-order conditions and use the restrictions that $\sum_s \pi_s = 1$ as well as $\sum p_s = 1$.

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:

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

That is, in equilibrium there exists a representative agent whose utility can be chosen to be quadratic as well.¹⁸ Moreover, the larger W_s^0 , the smaller the ratio p_s^*/π_s .

It follows from the discussion above 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.

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). 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).¹⁹

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 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

¹⁸We 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.

¹⁹The 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.

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.²⁰ 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.

On the other hand, 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 as a next step to 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 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.) To make our case starker, we shall further assume that all agents have *identical quadratic utilities*, with coefficients $a_n = a$ and $b_n = B$. Thus, any effect on equilibrium pricing *vis a vis* the analysis in the previous section will be due to the heterogeneity in ambiguity attitudes. Clearly, these results extend *a fortiori* if we allow utility functions to be nonquadratic (but see footnote 13) and/or heterogeneity in risk attitudes.

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, order them so that subjects $n = 1, \dots, N^*$ ($N^* < N$) have expected utility preferences and subjects $n > N^*$ have α -maxmin preferences with $\alpha_n = 1$ —i.e., they are extremely ambiguity averse. The equilibrium condition for the market of the pure-risk state security r is standard:

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

²⁰This 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.

(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:

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

where we used the fact that $w_g^{n,*} = w_b^{n,*}$ for $n > N^*$.²¹

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

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

The equilibrium state price probability ratios of the ambiguous state securities are more complicated. It is most useful to rewrite them as solutions to the following fixed point problem.

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

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

$$\begin{aligned}\frac{p_r^*}{\pi} &= \frac{a}{a - BE[W^0]} - \frac{B}{a - BE[W^0]} W_r^0; \\ \frac{p_g^*}{\rho} &= \left(\frac{a}{a - BE[W^0]} + \frac{\frac{1}{N} \sum_{n=N^*+1}^N \lambda^{n,*} p_g^* - p_b^*}{a - BE[W^0]} \frac{1}{1-\pi} \right) - \frac{B}{a - BE[W^0]} W_g^0; \\ \frac{p_b^*}{1-\pi-\rho} &= \left(\frac{a}{a - BE[W^0]} - \frac{\frac{1}{N} \sum_{n=N^*+1}^N \lambda^{n,*} p_g^* - p_b^*}{a - BE[W^0]} \frac{1}{1-\pi} \right) - \frac{B}{a - BE[W^0]} W_b^0.\end{aligned}$$

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 approximately 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 can be approximated by straight lines with different intercepts but common slope (the approximation being due to the fact that the difference $p_g^* - p_b^*$ is not constant, but we treat it as such). 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

²¹To 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.

states and that for the pure-risk state are the same. So, the equilibrium conditions for the ambiguous securities can be approximately 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 in the presence of specific patterns of the aggregate supplies of securities.

To understand this effect in more detail, consider *first* the case in which $W_g^0 < W_r^0 < W_b^0$. That is, the aggregate 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 $p_g^* > p_b^*$. *Intuitively*, what will happen is that the ambiguity averse agents will try to hedge the ambiguity by holding the two ambiguous securities in equal proportions, leaving the expected utility maximizers to absorb the supply imbalance. In order for the latter to be willing to do so, the price differential $p_g^* - p_b^*$ has to be larger than the one that would be observed in the absence of ambiguity averters. And indeed, inspection of the three equilibrium conditions 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 under these supply conditions be in accordance to theoretical predictions if there was no ambiguity*: 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 higher supply. *Perverse ranking* of state price probability ratios now becomes possible —i.e., perverse relative to predictions

ignoring ambiguity aversion. The intuition is analogous to the one given above: The difference $p_g^* - p_b^*$ has to be larger than the one that would be observed in the absence of ambiguity averters. The result may be that the state price probability ratio p_g^*/ρ ends being higher than p_r^*/π , contradicting the “standard” price ranking. 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* (relative to a model that ignores the presence of ambiguity averse subjects) becomes possible. The utility of the representative agent has to be made state-dependent, as in the second case above.

As mentioned earlier, nothing of what we described changes qualitatively if we allow heterogeneity in risk attitudes as well. However, it can be seen by re-elaborating our quadratic example that if there is *correlation between risk aversion and ambiguity aversion* —in the sense that subjects $n = N^* + 1, \dots, N$ 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 tend to 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. 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, $W_X^0 = 5 < W_Z^0 = 10 < W_Y^0 = 15$ in each experiment.

4.1 State Price Probability Ratios

In experiments 030123, 030203 and 030210, state securities X and Y, the securities in lowest and highest supply, were ambiguous. 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 3.2.1) or heterogeneous expected utility (with a uniform prior) and maxmin (Section 3.2.2).

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 some subjects did exhibit aversion against ambiguity.²² 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 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.

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, we first used probabilities that were updated from the 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.

As benchmarks to evaluate the findings in the two sets of experiments with ambiguous states, Figure 7 provides empirical distribution functions of the state-price probability ratios in the two pure-risk experiments each of which replicates a particular experiment with ambiguity. Experiment 040908 replicates 030203; experiment 041007 replicates 020529. These experiments are replications in the sense that the same per-capita supplies were imposed, exactly the

²²Under 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.

same draws were used, and approximately the same number of subjects (but from a different cohort) participated. The subjects, however, were informed of the exact composition of the urn, however. Both pure-risk experiments replicate findings in earlier work on asset pricing in an experimental setting [see ?]: state securities in low supply (X) tend to be expensive relative to their probability; states in high supply (Y) tend to be cheap.

The comparison between experiments 041007 (Figure 7) and 020529 (Figure 5) is especially revealing. In the latter, violations of the ranking expected under pure risk is highly likely; in the former, such ranking violations are unlikely.

We deflated state prices with probabilities based on the “obvious” 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. For instance, suppose that the marginal investor for the ambiguous securities is α -maxmin with $\alpha = 0.5$. As we argued in Section 3.1.2, 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, we used probabilities that were updated from the uniform prior. We now study a distinct 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 8 displays the results. A different 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 securities. 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.

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.

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. Figure 9 compares the histograms of final

holdings in experiment 040908 (pure risk; top) and in experiment 030203 (ambiguity; bottom); Figure 10 compares the histograms of final holdings in experiment 041007 (pure risk; top) and in experiment 020529 (ambiguity; bottom). The histograms show the frequency distributions of the proportion of end-of-period wealth exposed to ambiguity (if ambiguity is present) allocated to the ambiguous security in highest supply (Y) (under pure risk, the corresponding proportion is shown). In Figure 9, the frequency distributions of the ratio of number of securities Y over the total holdings of X and Y are displayed; In Figure 10, the frequency distributions of the ratio of number of securities Y over the total holdings of Y and Z are displayed.

In both comparisons, two effects stand out.

1. Under ambiguity (bottom histograms), the histograms are skewed to the right, with a peak around 0.5, where subjects are not exposed to ambiguity. Under pure risk, the histograms are slightly negatively skewed relative to the average per-capita supply ratios, which are approximately 0.75 (040907 and 030203) and 0.60 (041007 and 020529).
2. The cross-sectional variance of the proportions are far higher under ambiguity (030203 and 020529) than under pure risk (040907 and 041007). It is obvious why: the average per-capita supply ratios are large than 0.5, so that some (ambiguity tolerant) subjects need to hold far more than the average per-capita supply ratios in order to offset the ambiguity averse subjects who choose not to expose themselves to ambiguity.

Both effects are also evident in Figures 11 and 12, which plot the proportions (vertical axis) against the volatility of end-of-period wealth (horizontal axis). 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.²³

The figures also suggest a positive correlation between ambiguity aversion and risk aversion. Subjects that choose to hold ambiguous securities in less balanced proportion tend to also be exposed to more risk.

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

²³Volatility is measured in francs.

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

Table 1: Experiment Parameters

Date	Urn			Subject Category (Number)	Signup Reward (franc)	Endowments ^a			Loan (franc)	Exchange Rate cents/franc
	(18 Total)					X	Y	Z		
	X	Y	Z							
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
040907 ^b	6	3	9	15	250	4	11	4	500	2
				14	250	1	4	6	375	2
041007 ^c	6	6	6	15	500	4	11	2	220	1
				14	500	1	4	8	300	1

^aNumber of each of the state securities, in addition to 200 francs cash. No net holdings of Notes.

^bReplication of 030203 with the same draws but without ambiguity.

^cReplication of 020529 with the same draws but without ambiguity.




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).

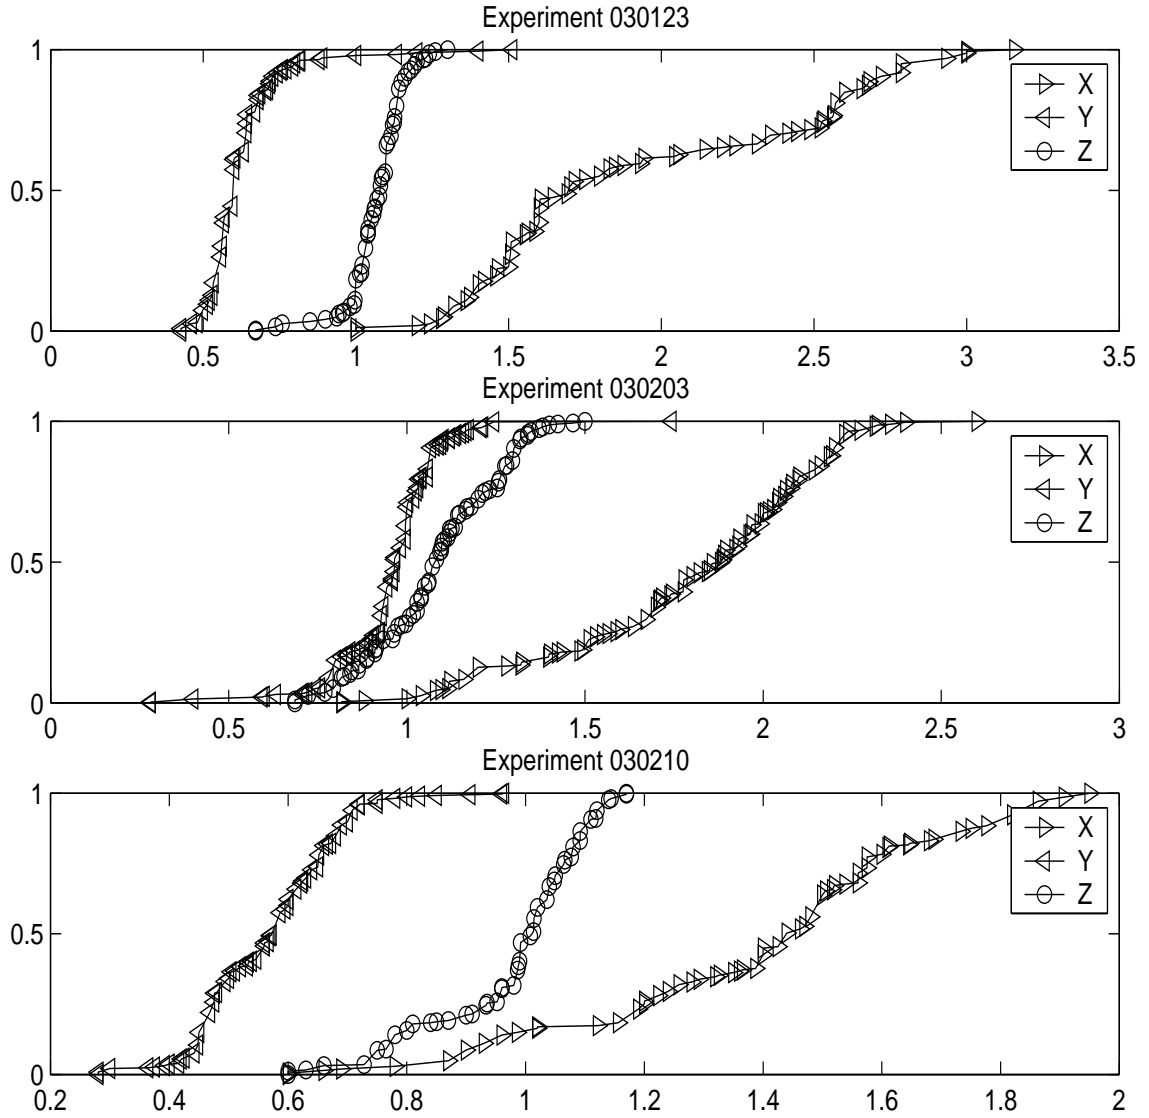


Figure 5: Empirical distribution functions of state price probability ratios in three experiments where the ambiguous securities, X and Y, are in lowest and highest supply, respectively. State price probability ratios are obtained by deflating state prices with probabilities based on (i) a diffuse prior 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).

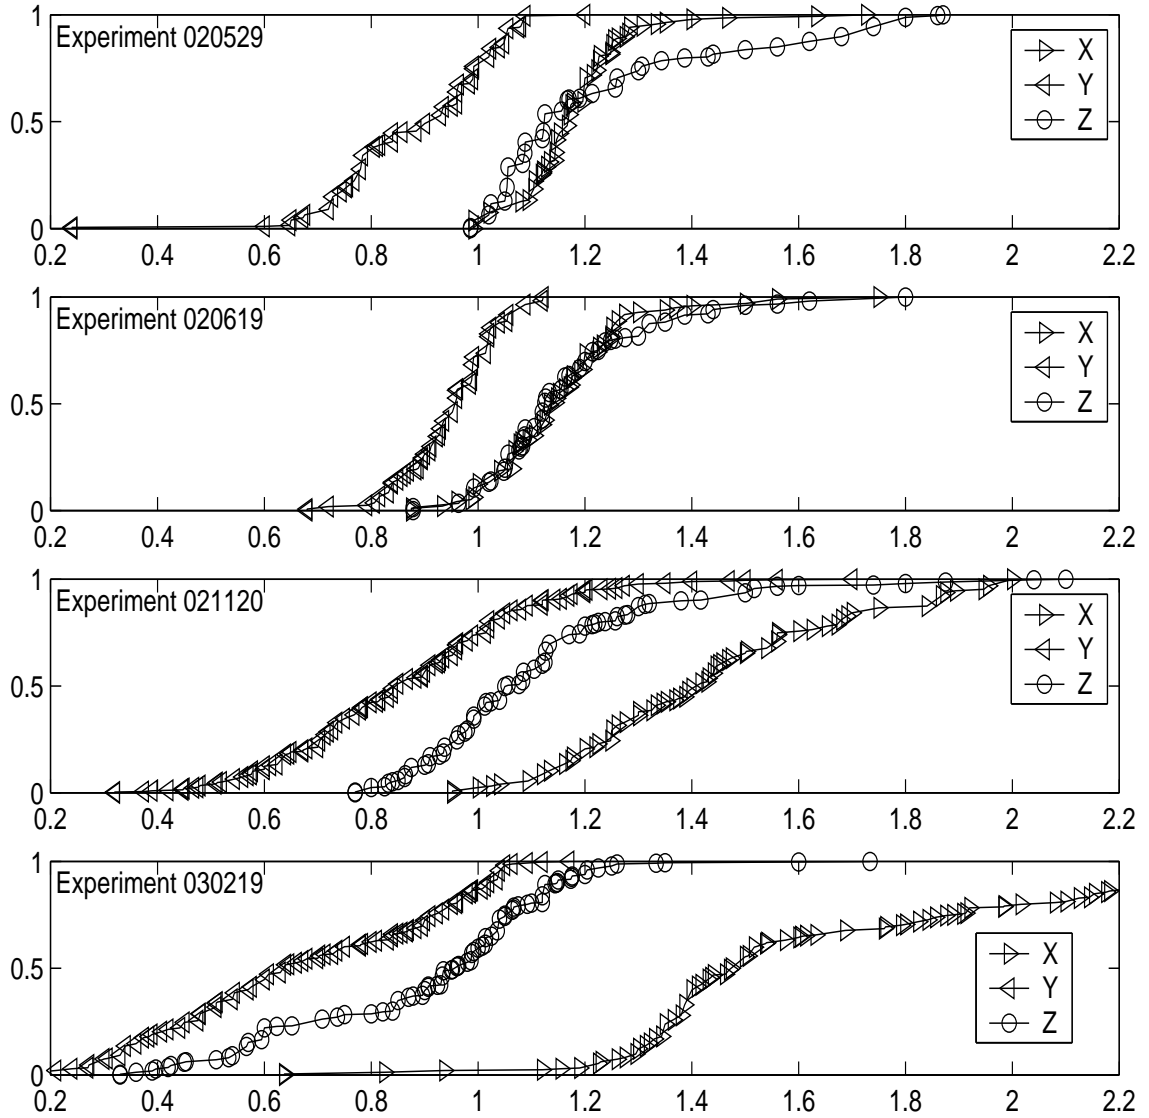


Figure 6: Empirical distribution functions of state price probability ratios in four experiments where the ambiguous securities, Y and Z, are in highest supply. (The pure risk state security, X, is in lowest supply.) State price probability ratios are obtained by deflating state prices with probabilities 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.

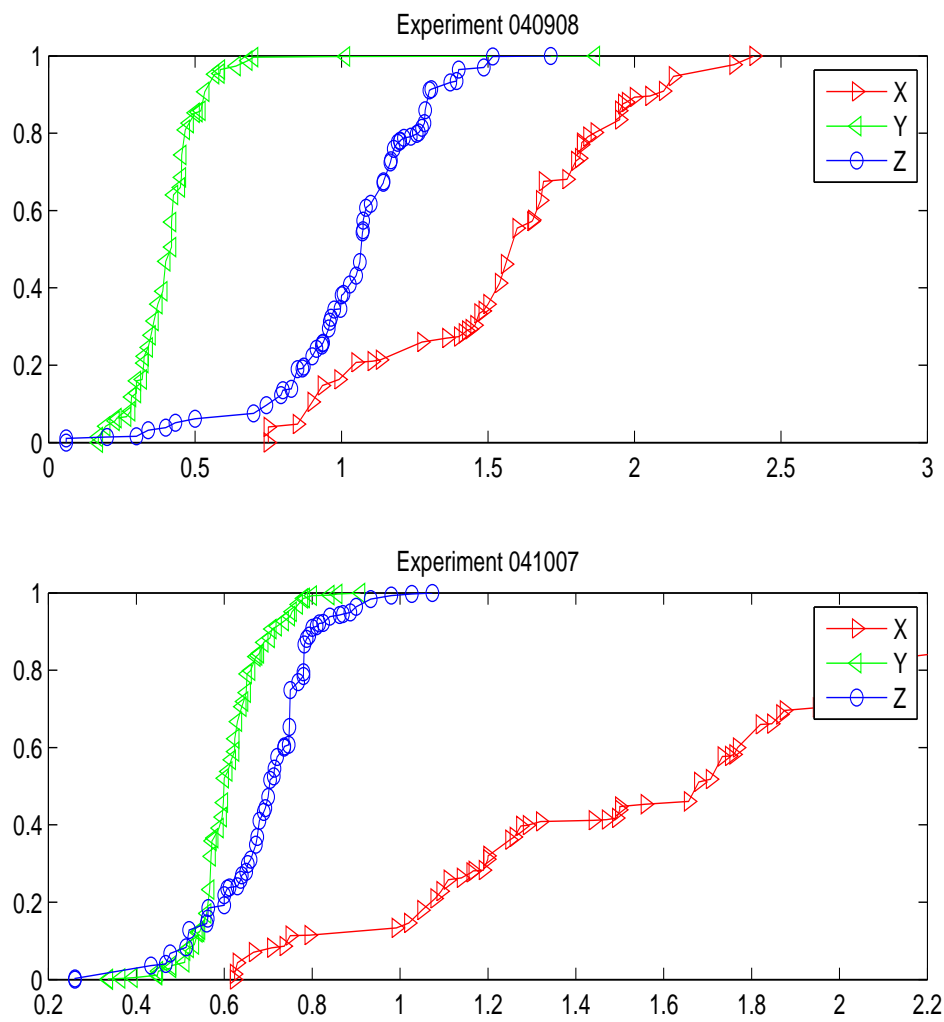


Figure 7: Empirical distribution functions of state price probability ratios in two experiments without ambiguity. Experiment 040908 is a replication of 030203 (same per-capita supplies of risky securities; same draws); experiment 041007 is a replication of 020529. Ordering of state-price probability ratios is as expected: X is most expensive, Y is cheapest.

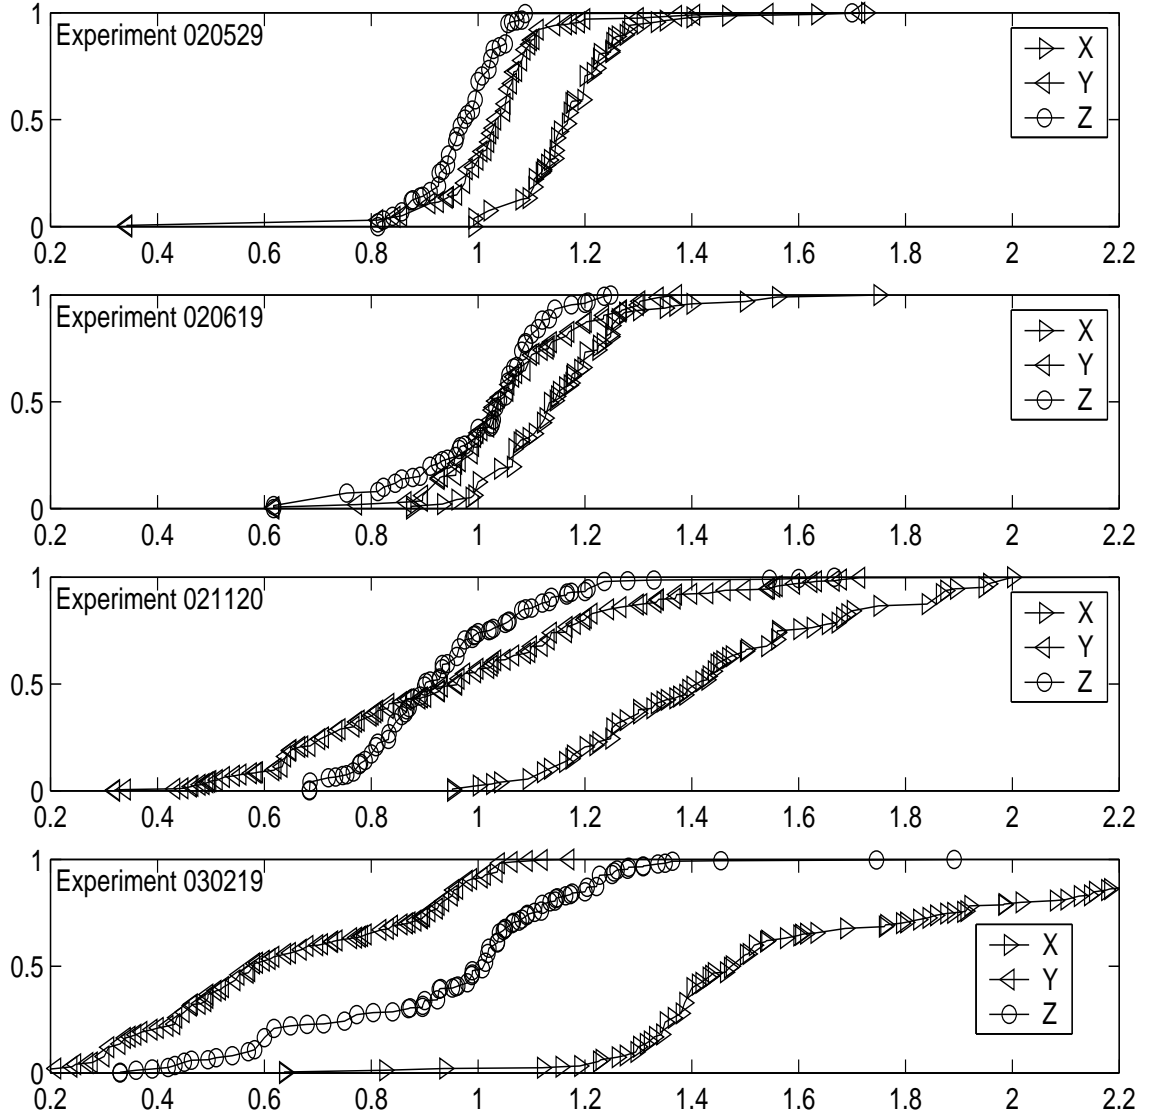


Figure 8: Empirical distribution functions of state price probability ratios in four experiments where the ambiguous securities, Y and Z, are in highest supply. (The pure risk state security, X, is in lowest supply.) State price probability ratios are obtained by deflating state prices with probabilities 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 9: Histograms of holdings of state security Y (in highest supply) as a proportion of holdings of state securities X (in lowest supply) and Y, in experiment 040908 (where no state was ambiguous; top) and in experiment 030203 (where states X and Y were ambiguous; bottom). Note the mode at 0.5 under ambiguity, and the higher cross-sectional variance.




Figure 10: Histograms of holdings of state security Y (in highest supply) as a proportion of holdings of state securities Z (in lower supply) and Y, in experiment 041007 (where no state was ambiguous; top) and in experiment 020529 (where states Y and Z were ambiguous; bottom). Note the mode at 0.5 under ambiguity, and the higher cross-sectional variance.




Figure 11: Plot of holdings of state security Y (in highest supply) as a proportion of holdings of state securities X (in lowest supply) and Y, in experiment 040908 (where no state was ambiguous; top) and in experiment 030203 (where states X and Y were ambiguous; bottom) against volatility (in francs) of end-of-period individual wealth (horizontal axis). Note the positive correlation between risk aversion and exposure to ambiguity in experiment 030203.

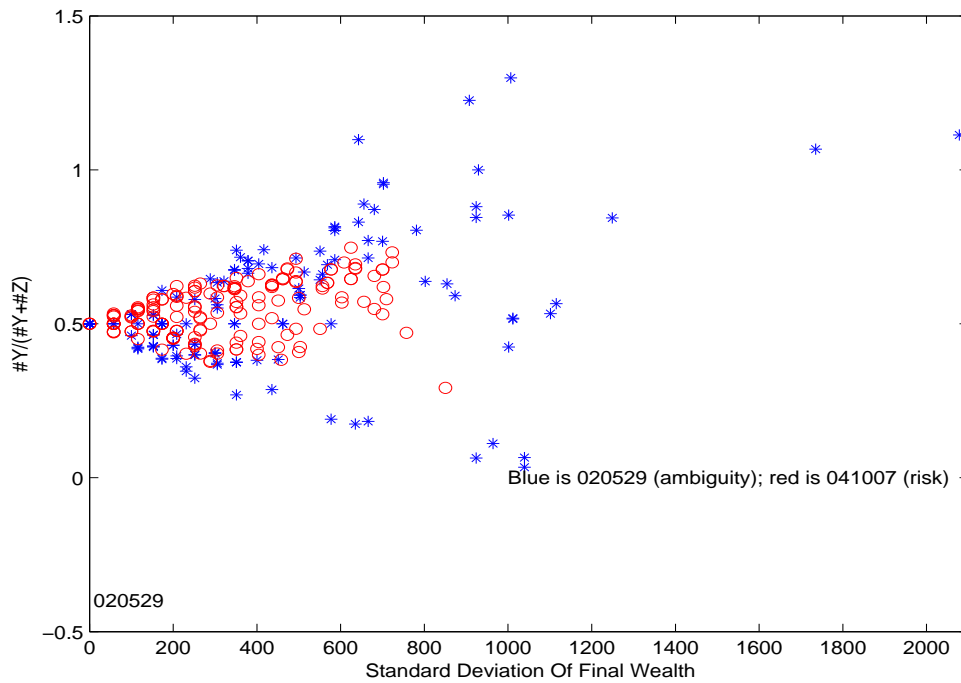


Figure 12: Plot of holdings of state security Y (in highest supply) as a proportion of holdings of state securities Z (in lower supply) and Y, in experiment 041007 (where no state was ambiguous; top) and in experiment 030203 (where states Y and Z were ambiguous; bottom) against volatility (in francs) of end-of-period individual wealth (horizontal axis). Note the positive correlation between risk aversion and exposure to ambiguity in experiment 020529.