



Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets*

PETER BOSSAERTS¹ and CHARLES PLOTT²

¹*California Institute of Technology and CEPR;* ²*California Institute of Technology*

Abstract. We report on two sets of large-scale financial markets experiments that were designed to test the central proposition of modern asset pricing theory, namely, that risk premia are solely determined by covariance with aggregate risk. We analyze the pricing within the framework suggested by two theoretical models, namely, the (general) Arrow and Debreu's complete-markets model, and the (more specific) Sharpe-Lintner-Mossin Capital Asset Pricing Model (CAPM). Completeness of the asset payoff structure justifies the former; the small (albeit non-negligible) risks justifies the latter. We observe swift convergence towards price patterns predicted in the Arrow and Debreu and CAPM models. This observation is significant, because subjects always lack the information to deliberately set asset prices using either model. In the first set of experiments, however, equilibration is not always robust, with markets temporarily veering away. We conjecture that this reflects our failure to control subjects' beliefs about the temporal independence of the payouts. Confirming this conjecture, the anomaly disappears in a second set of experiments, where states were drawn without replacement. We formally test whether CAPM and Arrow–Debreu equilibrium can be used to predict price movements in our experiments and confirm the hypothesis. When multiplying the subject payout tenfold (in real terms), to US \$ 500 on average for a 3-h experiment, the results are unaltered, except for an increase in the recorded risk premia.

* Financial support was provided by the Caltech Laboratory of Experimental Economics and Political Science, the National Science Foundation, the International Center For Finance at Yale University and a grant from the Jenkins Family Foundation to Caltech. We thank Bill Brown (Claremont), Peter DeMarzo (Berkeley, now at Stanford), Will Goetzmann (Yale), Mark Johnson (Tulane), Tony Kwasnica (Penn State), Claude Montmarquette and Claudia Keser (CIRANO, Montréal), William Sharpe (Stanford), Woody Studenmund (Occidental College), and Ivo Welch (UCLA, now at Yale), for allowing us to involve their students in our experimental financial markets. Elena Asparouhova (now at Utah) facilitated the Bulgarian experiment. The paper benefited from comments at many seminars and meetings throughout the world. Constructive criticism from the editor and two referees is gratefully accepted, but only the authors are responsible for remaining errors. Correspondence to: pbs@rioja.caltech.edu, cplott@hss.caltech.edu.

1. Motivation

Since the early 1950s, two major considerations have driven the development of asset pricing theory, namely, (i) the presence of risk-averse investors, and (ii) asymmetric information. The former consideration has led to a better theoretical understanding of risk and its equilibrium compensation. The predictions from specific versions of the theory are widely applied in practical areas, such as capital budgeting and portfolio performance evaluation, yet empirical support from econometric analysis of historical data has been dubious at best. The disconnect between theory and empirical support calls for an experimental investigation, which allows for better control of many of the key assumptions. Asymmetric information adds complexity to an already controversial theory. Also, with few exceptions (e.g., Biais, et al. 2003), the literature on competitive asset pricing under asymmetric information is entirely theoretical. (There is a large literature of pricing under asymmetric information in a *strategic* setting, though.) Therefore, we decided to avoid asymmetric information in the first stage of our experimental investigation.

The premises of standard asset pricing theory have been that expected utility maximizing agents meet in competitive markets, and that prices react, to eventually reach general equilibrium. The main implications have been that equilibrium prices will only reflect compensation for aggregate, i.e., non-diversifiable, risk, and that this risk is to be measured in terms of the covariance between the return on a security and some aggregate risk factor(s). Preferably, the latter can be identified in terms of the return on a specific (set of) benchmark portfolio(s).

While competitive market models are a fiction, they are capable of making accurate predictions in the laboratory if one uses the right trading interface and if a sufficient number of subjects are present. Several choices are available as far as trading interface is concerned. In our experiments, we use a continuous, web-based (electronic) open book system similar to the one in place in many stock markets around the world. In such a setting, however, earlier experiments (Bossaerts and Plott 2002) did not produce the expected result: prices did not converge to equilibrium configurations. Bossaerts and Plott (2002) conjectured that the number of subjects (5–13), while standard in experimental economics, is insufficient to make markets liquid enough in a portfolio setting.¹ Their reasoning was as follows. Asset pricing theory becomes non-trivial only when there are at least three securities. As part of portfolio rebalancing, subjects may want to sell one security to buy another

¹ Experimentalists consider 8–12 subjects to be sufficient to generate convergence towards competitive equilibrium in the laboratory. These numbers come from much simpler experiments, however, where at most two goods/securities are traded.

one. Since all trading takes place through cash (swaps are not possible), subjects may end up with an inferior portfolio if, because of market thinness, they cannot complete portfolio rebalancing before the end of the period. Subjects may therefore refrain from rebalancing. As a result, prices do not converge to equilibrium. One of the goals of this paper is to verify whether this reasoning is correct, by increasing the number of subjects substantially. In our large-scale experiments, the number of subjects varied between 19 and 63.²

Asset pricing theory builds on risk aversion. Provided they have smooth expected utility preferences, subjects should be approximately risk neutral when risk is small. Therefore, one wonders whether risk in laboratory experiments can be raised enough for significant risk premia to emerge. There is, however, plenty of evidence in the experimental economics literature that significant risk aversion emerges at the levels of uncertainty present in our experiment, in the form of low transaction prices relative to expected payoffs.³ But note that there are other theories besides risk aversion that could explain discounts, such as the endowment effect. Our study casts new light on the origin of the discounts. Asset pricing theory not only predicts that risk premia will emerge, but also how these are to be distributed across assets. So, when the cross-section of the discounts becomes patterned after the theory, as it does in our experiments, the hypothesis that risk aversion is at work gains credibility.

One way to be convinced that the observed discounts are really risk premia is to increase the size of the risk. In one experiment, we effectively increased risk (and return) at least 10-fold. Subjects were Bulgarian students and we offered them nominally what we had been paying North American subjects. Since the standard of living of the Bulgarian students is at most one-tenth that of the North American students (at the time of the experiment), the real-term compensation and risk were at least 10 times as high for the Bulgarians.

In field studies, static asset pricing models are often used to interpret single-period returns on multi-period securities. Our experiments are closer to the theory, because our securities are single-period. In particular, our experiments are sequences of replications of the same situation; a replication is referred to as a period. Neither cash nor asset holdings are carried over from one period to another. Still, to mitigate risk-loving behavior because of limited liability (in the worst case, subjects leave the experiment empty-handed), we introduced a “solvency rule.” The rule creates subtle links

² As an alternative, one could explore more sophisticated trading mechanisms, meant to alleviate the portfolio rebalancing problem, such as the combined-value auction systems discussed in Bossaerts, Fine and Ledyard (2002). The need to increase the number of subjects when studying more complicated general equilibrium is not unique to asset pricing. See, e.g., Noussair, Plott and Riezman (1997).

³ See Holt and Laury (2002) for an extensive investigation of the phenomenon.

between periods. It works as follows. Subjects' earnings cumulate across periods; negative earnings in a period are to be offset with positive earnings in other periods; subjects with negative cumulative earnings for two periods in a row are declared insolvent and barred from trading in future periods. A subject who is insolvent foregoes an opportunity to gain in the future. This effectively induces caution in all periods but the last.⁴ Our solvency rule therefore generates subtle dynamic features not captured by the static asset pricing models we use to interpret the data. Still, these dynamic features had no noticeable effect on pricing: risk premia do not decrease over time. (Nor did they have significant effects on choices: end-of-period holdings of the vast majority of subjects are rationalizable in terms of a fixed, concave utility function.)

Asset pricing theory makes predictions about equilibrium prices that depend on agents' preferences and beliefs. In the case of the model of Arrow and Debreu, strong ordinal predictions about ratios of state prices over probabilities obtain irrespective of individual preferences, as long as utility is smooth and beliefs are common. CAPM makes cardinal predictions about the pricing of one portfolio, namely, the market portfolio, but imposes stronger requirements on utility: it is assumed to be quadratic. Smooth utility is only mildly controversial.⁵ Even quadratic utility can be justified, at least in our experimental setting, because risk is small, so that quadratic utility obtains as an approximation of true utility. It may seem that commonality of beliefs is a not an issue, given clear instructions that explain how states (which determine payoffs) are drawn. As we shall see, however, prices at times reflect a disbelief in the temporal independence of the drawing across

⁴ Consider the penultimate period. If a subject manages to satisfy the solvency rule in that period, she receives securities with which to make money in the last period. To maximize the probability of receiving these valuable securities, our subject has to minimize the probability of earning less than is required to satisfy the solvency rule. If the payoff distribution is symmetric, maximizing variance, as a risk-neutral subject would do if the penultimate period were really the last one, maximizes this probability instead of minimizing it (matters are more complicated if the payoff distribution is asymmetric). Consequently, the subject has to balance maximizing expected payoff in the penultimate period against the probability of being thrown out. Incidentally, maximizing payoff variance is wrong for a second reason. The payoff in the penultimate period indirectly determines the nature of the securities in the last period. Indeed, because we cannot force subjects to pay for losses, the securities they hold actually form a combination of full-liability securities plus a put option. The put option allows her to walk away from the experiment if she ends with a negative balance. The strike price of the put option is determined by the cumulative earnings up to the penultimate period. The higher these cumulative earnings, the lower the strike price of the put option. To increase the value of the put option, our subject has to increase the strike price, and hence, to reduce the chances of getting high earnings in the penultimate period. Again, maximizing variance, as a risk-neutral subject would do if the penultimate period were the last one, is not optimal.

⁵ Recently, asset pricing models have been suggested that build on kinks in the utility function at endowment. See, e.g., Barberis, Huang and Santos (2001), or Epstein and Zin (2001).

periods. We conjectured that subjects have a natural inclination to better understand (and believe) draws without replacement. To verify this, we ran an additional three experiments where states are drawn without replacement. The latter will be referred to as the second set of experiments, to distinguish them from the first set, which consists of six experiments where states were drawn with replacement.

The significance of our experimental results is further enhanced by the following design feature. Subjects knew only their own endowment and not that of others.⁶ So, they did not know the aggregate endowment. Yet, to price securities in accordance to the Arrow–Debreu model or to the CAPM, the aggregate endowment has to be known. For instance, the CAPM predicts that prices will be set such that the market portfolio (i.e., the aggregate portfolio of risky securities) is mean-variance optimal. Subjects did not know the market portfolio, and hence, could not use the CAPM to price securities. If CAPM pricing emerges in our experiments, it cannot be attributed to subjects' using CAPM to determine offer prices.

We hasten to add that *neither the Arrow–Debreu model nor the CAPM require agents to know the aggregate endowment for its pricing predictions to be valid*. Prices in both models are *Walrasian equilibrium prices*: agents see prices and optimize; the prices are such that their optimal demands equal supplies. Absent asymmetric information,⁷ agents need not know anything about other agents' preferences, beliefs, endowments, or even the aggregate supply of securities (the market portfolio).⁸

Very little experimental work on symmetric-information static asset pricing has been done. There is an older literature on dynamic asset pricing (bubbles), but risk neutrality is invariably assumed. We know of no work on the Arrow–Debreu equilibrium. A few studies have focused on the CAPM. Kroll, Levy and Rapaport (1988) mainly looked at individual portfolio choice (in particular, whether individual holdings could be separated into a position in the riskfree security and a position in the market portfolio). Levy (1997) studied the CAPM in a repeated call market with 20 securities and discovered a significantly positive relationship between mean excess return and covariance with the market return, providing partial support for the

⁶ They were not even told how many subjects participated and could not determine the number because our experiments were web-based, and therefore, decentralized.

⁷ Under asymmetric information, knowledge of the aggregate endowment does become important, at least if one entertains a notion of equilibrium that is different from the Walrasian equilibrium, such as the rational expectations equilibrium. See, e.g., Biais, Bossaerts and Spatt (2003).

⁸ From an empirical point of view, however, our experiments provide a solution to the Roll critique. Roll (1977) questioned tests of the CAPM on field data because the empiricist needs to observe the market portfolio and cannot substitute even a highly correlated proxy. In the experiments, we are in control of the market portfolio, and hence, know it.

CAPM. (As mentioned before, the CAPM actually imposes a much stronger restriction, namely, that this relationship be proportional.) Bossaerts and Plott (2002) discusses results of a set of experiments based on a design that is almost identical to the one described here. Far fewer subjects participated, however, thus reducing competition and ultimately inhibiting convergence, as discussed earlier.

Formal inference is hampered by the absence of a suitable measure of absolute distance from equilibrium. It is difficult to determine when markets are close to equilibrium in a theory that focuses exclusively on equilibrium. As a consequence, our statistical tests will not examine to what extent markets are close to equilibrium at a given point, but, more modestly, whether one can use equilibrium theory to predict in what direction prices move.

The remainder of this paper is organized as follows. Section 2 describes the experimental setup. Section 3 spells out the theoretical predictions. Evidence is reported in Section 4. Section 5 concludes.

2. Experimental Design

The experiments are organized as follows. Three assets can be traded. Two, labeled Securities A and B, have a risky payoff, determined by the random drawing of one of three “states,” referred to as States X, Y and Z. The third asset, labeled Note, is riskfree, and, unlike the risky securities, can be sold short, up to eight units.⁹ At the beginning of each period, subjects are endowed with a certain number of Securities A and B and cash (purchases are paid for in cash). The Notes are in zero net supply.¹⁰

During 15–25 min, subjects can submit limit orders 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 in the sense that only abstract IDs (numbers between 100 and 200) are shown; 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

⁹ In neither the Arrow–Debreu equilibrium nor in the CAPM equilibrium are short-sale constraints on risky securities binding. However, the ability to short sell the riskfree security is potentially important.

¹⁰ Cash is a means of exchange and an asset alike. Notes are only an asset (i.e., cannot be used to buy other securities). As assets, cash and notes are perfect substitutes. But cash provides a service (payment) which may cause it to be priced differently, at least temporarily (as long as it provides a payment service). We elaborate when we discuss the empirical results later on.

discussed below. Subsequently, assets are taken away, and a new period starts (subjects are given a fresh supply of the risky securities, as well as cash, etc.). Each experiment involved up to eight periods.

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 Securities A and B. 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. As mentioned in Section 1, this *solvency rule* induces risk-averse behavior even among risk-neutral subjects, except in the last period.

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 A	170	370	150
Security B	160	190	250
Notes	100	100	100

The remaining data and parameters for the experiments are displayed in Table I. There is some variation in initial endowments across experiments, for two reasons. First, the pricing implications ought to be robust to changes in endowments (provided, of course, that endowments do not reduce competition). For example, in the CAPM, the market portfolio is mean–variance optimal whether agents are initially endowed with the market portfolio (as in experiment 981007) or not. Second, while most subjects were inexperienced, a small minority of Caltech students participated in two experiments (at most). We wanted to avoid that these repeat-subjects used their experience in a previous session to better predict where prices were going.

When subjects are given different endowments, the parameters (including the loan repayments) are set such that subjects had similar risk/return tradeoff if they trade up to mean–variance optimal portfolios at one of the

Table I. Experimental parameters

Date	Draw type ^a	Subject category (number)	Signup reward (franc)	Initial		allocations notes	Cash (franc)	Loan (franc)	Exchange rate(\$/franc)
				A	B				
981007	I	30	0	4	4	0	400	1900	0.03
981116	I	23	0	5	4	0	400	2000	0.03
		21	0	2	7	0	400	2000	0.03
		8	0	5	4	0	400	2000	0.03
990211	I	11	0	2	7	0	400	2000	0.03
		22	175	9	1	0	400	2500	0.03
990407	I	22	175	1	9	0	400	2400	0.04
		33	175	5	4	0	400	2200	0.04
991110	I	30	175	2	8	0	400	2310	0.04
		22	175	5	4	0	400	2200	0.04
991111	I	23	175	2	8	0	400	2310	0.04
		21	125	5	4	0	400	2200	0.04
011114	D	12	125	2	8	0	400	2310	0.04
		18	125	5	4	0	400	2200	0.04
01126	D	18	125	2	8	0	400	2310	0.04
		17	125	5	4	0	400	2200	0.04
011205	D	17	125	2	8	0	400	2310	0.04
		17	125	2	8	0	400	2310	0.04

^aI = States are drawn independently across periods; D = states are drawn without replacement, starting from a population with 18 balls, six of each type (state).

possible equilibrium price sets (i.e., a set of prices such that the market portfolio had the highest possible Sharpe ratio). Also, the loan repayments are set such that a subject wishing to unload all risky securities would make some money, even if minimal.

The actual payoffs obviously depend on the sequence of states that happened to be drawn during the experiment. In the first six experiments, to be referred to as the *first set*, states are drawn with replacement, implying that the probability of each state does not change over time. In the last three experiments, to be referred to as the *second set*, states are drawn without replacement.¹¹ As mentioned in Section 1, price dynamics in some of the

¹¹ We start the experiment with an urn containing six balls labeled X, six balls labeled Y, and six balls labeled Z. At the end of the first period, a ball is drawn, determining the state for that period. This ball is not replaced, thus changing the relative likelihood of the states in subsequent periods. Subjects were fully informed about this procedure, through the instructions, and through occasional announcements during the experiment.

experiments in the first set (in particular, 990407 and 991111) indicated that subjects did not understand or did not believe that states were drawn independently. Instead, prices behaved as if they believed that the likelihood of a state decreased after it was drawn. To ensure better control of beliefs, we decided to switch protocols and draw states without replacement in the second set of experiments. Note, however, that each period in the second set, markets have to search for an equilibrium with *new* state probabilities. The type of drawing (with, without replacement) for each experiment is indicated in the second column of Table I.

We had limited control over the number of participants in the experiments. Subjects had to register in a database several days before the actual experiment. After that, they could retrieve an ID and password with which to log on to the website from which the experiment would be run, and execute some trades as practice. In general, however, subjects were slow to register. Hence, the number of subjects was random. In one experiment, we only had 19 subjects (the one referred to as 990211). At the same time, this gives us the opportunity to study the effect of number (of agents) on distance from equilibrium pricing across experiments.

Subjects are given instructions that they can read on the website before each experiment. The instructions are the same throughout, except for the description of the drawing of the states, which is different in the second set of experiments. In experiment 011126, subjects were Bulgarian, and a Bulgarian translation of the instructions was available. (Except for bilingual announcements, everything during this experiment was in English.)

It should be pointed out that most subjects had familiarity with financial markets. This is obvious for MBA students, which constituted the majority of the subjects in the first six experiments.¹² Caltech students were a minority from the second experiment on, and a majority in the last two experiments. All of them had had exposure to finance and/or micro-economics classes. A number of economics students from Claremont and Occidental College, as well as the University of Montreal also participated in the later experiments. Only Bulgarians from the University of Sofia take part in experiment 011126. Few of these had had prior financial experience.

Average per-subject payments in our experiments is approximately US \$50; maximum pay is about US \$250 (in experiment 981116, which was a 4-h experiment, exceptionally); minimum pay is US \$0.

The websites for the experiments are as follows:

<http://eeps2.caltech.edu/market-981007yale> (“981007”)

¹² In 981007 and 990211, students were MBAs from Yale; in 981116, they came from UCLA (MBAs); in 990407, Stanford MBAs were used; in 991110, Tulane University MBAs participated; in 991111, the majority of the subjects were Berkeley MBAs.

<http://eeps2.caltech.edu/market-981116uclacit> (“981116”)

<http://eeps2.caltech.edu/market-990211> (“990211”)

<http://eeps2.caltech.edu/market-990407> (“990407”)

<http://eeps2.caltech.edu/market-991110> (“991110”)

<http://eeps2.caltech.edu/market-991111> (“991111”)

<http://eeps2.caltech.edu/market-011114> (“011114”)

<http://eeps2.caltech.edu/market-011126> (“011126”)

<http://eeps2.caltech.edu/market-011205> (“011205”)

The interested reader can visit these websites, read the instructions, inspect the trading interface, and display the trading history (including pricing).¹³

3. Theoretical Predictions

At the core of modern asset pricing theory is the hypothesis that competitive financial markets *equilibrate* and that, when markets are complete and agents are expected utility maximizers, equilibrium risk premia can be characterized in terms of the *covariation of an asset’s return with aggregate risk*. Complete markets is the hallmark of the Arrow–Debreu economy: all risks can be insured. (Debreu, 1959). Such is the case in our experiments. In addition, our experiments consist of non-overlapping periods, so we are using a static model to generate pricing predictions.¹⁴

Let R_i denote the return (payoff divided by price) on asset i and let $u'(W)$ denote the marginal utility u of wealth W of some “aggregate” investor, which exists by virtue of market completeness. We implicitly take the aggregate investor’s utility to be smooth, so that marginal utilities exist. We assume that $u'(W)$ is strictly positive (no satiation) and decreasing in wealth (which means that the aggregate investor is risk averse). From the first-order conditions for optimality of the aggregate investor’s investment–consumption plan, it must be that:

$$E[u'(W)R_i] = \lambda$$

¹³ Anonymous login requires the ID 1 and password a.

¹⁴ Still, our solvency rule induces subtle dynamic features – but which do not show up in the pricing, as we report later on.

for some $\lambda > 0$.¹⁵ Let R_F denote the return on the riskfree security. The above restricts R_F as well, so that we can restate the first-order conditions as follows:

$$E[R_i - R_F] = -\text{cov}\left(R_i, \frac{u'(W)}{E[u'(W)]}\right). \quad (1)$$

The task of asset pricing theory has been to identify the aggregate investor, so that $u'(W)$ can readily be measured. (Terminal) aggregate wealth equals the payoff on the aggregate portfolio of all investments in the economy, which has become known as the *market portfolio*.

It deserves emphasis that Equation (1) generates a set of equations that define the Walrasian equilibrium in an Arrow–Debreu world. That is, they produce a set of non-linear equations, to be solved for prices. The latter are implicit in Equation (1). To make them explicit, let X_i denote the payoff on asset i , let P_i denote its price, let P_F denote the price of the riskfree security, and assume that the latter's payoff is 1. Then:

$$E\left[\frac{X_i}{P_i} - \frac{1}{P_F}\right] = -\text{cov}\left(\frac{X_i}{P_i}, \frac{u'(W)}{E[u'(W)]}\right). \quad (2)$$

When markets equilibrate, they effectively solve the set of non-linear equations in (2).

In words, Equation (1) predicts that equilibrium securities prices will be such that risk premia (expected excess returns) are determined by covariation with aggregate risk, where aggregate risk is to be measured in terms of the marginal utility of an aggregate investor. This implies not only that risk premia will exist, but also what their form will be. Let us formally state the first part, and then study more closely the second part.

Risk Pricing Property: Risk premia will exist, i.e., equilibrium prices will generally deviate from expected payoffs.

The restrictions in Equation (1) describes equilibrium prices in an Arrow–Debreu world where agents have smooth expected utility preferences. They can be converted in restrictions on the prices of (imaginary) state securities that are easily verifiable in an experimental setting (where aggregate wealth is observable). To see this, let there be S possible states, indexed $s = 1, \dots, S$, each with probability π_s ($\sum_{s=1}^S \pi_s = 1$). Let P_s denote the price of state security s (that is, the security that pays US \$1 if state s occurs, and US \$ 0 otherwise). Let W_s denote the wealth of the aggregate investor in state s . Re-arranging Equation (1), one obtains:

¹⁵ In the experiments, periods are short, so that we expect $\lambda = 1$, i.e., the riskfree rate equals zero. We confirm this in the experiments (see the discussion of the empirical results). For the theory, we will leave λ (and the riskfree rate) unspecified, as our arguments do not rely on any specific value.

$$P_s = P_F \pi_s \frac{u'(W_s)}{E[u'(W)]}. \quad (3)$$

Therefore, two states s and s' with equal likelihood ($\pi_s = \pi_{s'}$) have prices that are ranked solely because of differences in aggregate wealth. If $W_s > W_{s'}$, then $u'(W_s) < u'(W_{s'})$ (because of decreasing marginal utility, i.e., risk aversion), and, hence, $P_s < P_{s'}$. The ranking of the state prices is the inverse of the ranking of aggregate consumption (wealth).

When two states have different likelihoods, then the rank order restriction does not apply to state prices, but to the ratio of state prices over probabilities, referred to as *state price probability ratios*, and denoted ρ_s :

$$\rho_s = P_s / \pi_s. \quad (4)$$

This observation is the foundation for another theoretical property that we want to study in experimental markets.

State Price Probability Ratio Property: The ranking of state price probability ratios will be the inverse of the ranking of aggregate wealth in those states.

The restriction in Equation (3) allows state prices to be *above* state probabilities. That is, (implicit) state securities may trade at a premium to expected payoff, implying that risk premia may be negative. We verify this seemingly counter-intuitive implication in the experimental data.

The State Price Probability Ratio Property is ordinal. To produce a cardinal prediction, we need to be more specific about the shape of the utility function of the representative agent. We have only been assuming that the utility function is smooth. Our experimental setting allows us to be more specific. While non-negligible, risk is small. This justifies a quadratic approximation of subjects' actual preferences. Under quadratic preferences, the Capital Asset Pricing Model (CAPM, developed in Sharpe 1964; Lintner 1965; Mossin 1966) obtains. In other words, the level of risk in our experiments gives the Arrow–Debreu model cardinal content because CAPM pricing will obtain.¹⁶

The CAPM predicts that risk premia will be proportional to the covariance between a security's return and the return on a specific benchmark

¹⁶ We emphasize that we employ quadratic utility only as a *local approximation*. As a global description of utility, it is highly unsatisfactory, because it violates the axiom of non-satiation on which the Arrow–Debreu model is built. That is, for sufficiently high wealth levels, subjects are satiated. As Dybvig and Ingersoll (1982) show, this conflicts with the non-satiation on which absence of arbitrage is based. In particular, they show that arbitrage opportunities become possible (if the market portfolio has sufficiently high payoff in some states). Later on, we will demonstrate directly that arbitrage opportunities almost never emerge in our experiments, because state securities prices are always strictly between zero and one and add up to one.

portfolio, namely, the market portfolio (i.e., the value-weighted portfolio of all risky securities). To see this mathematically, re-consider Equation (1). Utility is quadratic, and, hence, $u'(W) = a - bW$, for two positive coefficients a and b . Re-arranging Equation (1), one obtains:

$$E[R_i - R_F] = \frac{b}{E[u'(W)]} \text{cov}(R_i, W).$$

Let P_M denote the price of the market portfolio (the aggregate portfolio of investments in the economy). Its payoff, X_M , equals W . Let R_M denote its return, i.e., X_M/P_M . For the market portfolio,

$$E[R_M - R_F] = \frac{bP_M}{E[u'(W)]} \text{var}(R_M).$$

Using this to solve for $b/E[u'(W)]$, one obtains:

$$E[R_i - R_F] = \text{cov}(R_i, R_M) \frac{E[R_M - R_F]}{\text{var}(R_M)}. \quad (5)$$

In words, risk premia are proportional to covariances with the return on the market portfolio. The constant of proportionality is given by the expected excess return on the market portfolio divided by its variance.¹⁷

Roll (1977) shows that Equation (5) obtains whenever the market portfolio is mean–variance optimal, i.e., generates maximum expected return given its risk (return variance). This means that Equation (5) can be verified in a simple way using the *Sharpe ratio* (reward-to-risk) ratio of the market portfolio, i.e.,

$$\frac{E[R_M - R_F]}{\sqrt{\text{var}(R_M)}}.$$

¹⁷ Conventionally, the CAPM is re-written as follows:

$$E[R_i - R_F] = \beta_i E[R_M - R_F],$$

where

$$\beta_i = \frac{\text{cov}(R_i, R_M)}{\text{var}(R_M)},$$

i.e., the slope coefficient (“beta”) in an OLS projection of R_i onto R_M .

The market portfolio is mean–variance optimal if and only if the market's Sharpe ratio is maximal. Away from the CAPM equilibrium, the market's Sharpe ratio may not be maximal. Hence, the difference between the market's Sharpe ratio and the maximal Sharpe ratio provides a simple measure of how far the market is from the CAPM equilibrium. Thus, CAPM implies the following property.

Sharpe Ratio Property: The Sharpe ratio of the market will be maximal. Both the State Price Probability Ratio Property and the Sharpe Ratio Property have one important advantage. They provide a description of market equilibrium irrespective of the distribution of risk aversion or initial endowments across investors. Risk aversion and endowments may change (endowments definitely change once subjects trade), and hence, to have properties of equilibrium that do not require risk aversion or endowments to remain constant is a definite plus.

4. Results

The evidence on the theoretical properties will be presented here.

4.1. EVIDENCE FOR THE RISK PRICING PROPERTY

Figure 1 displays the evolution of the prices of the three securities in the first experiment, 981007. Each observation corresponds to a trade in one of the three securities (the price of the other two securities was set equal to their last trading price). The picture is representative for what happened in all the experiments, not just in 981007. Time is on the X -axis; prices are on the Y -axis. Vertical lines delineate periods; horizontal lines depict expected payoffs.

The first obvious observation one can make is that prices are always below expected payoffs (230, 200 and 100, for Securities A, B and Notes, respectively). This suggests support for risk pricing property, and reveals that risk aversion plays an important role in experimental financial markets, and hence, must be taken into account explicitly when assessing the success of the experiments. Such is not a foregone conclusion, however. Among other things, negative risk premia (in which case prices are above expected payoffs) can be consistent with risk aversion. What needs to be determined is whether the relative risk premia of A and B are as predicted by the theory. We will verify this by means of the State Price Probability Ratio and Sharpe Ratio Properties.

There is a striking cyclical in the prices of the Notes. Prices start out low in each period, but converge to 100 (the payoff), implying that the riskfree rate decreases steadily over the course of a period. Because there is no time value of money, the equilibrium interest rate must be zero. One generally observes near zero interest rates at the end of each period. The




Figure 1. Evolution of securities prices in the 981007 experiment. Vertical lines delineate periods; horizontal lines depict expected payoffs (top line: security A; middle line: security B; bottom line: Notes). All transactions of risky securities occurred at prices strictly below expected payoffs, suggesting risk aversion.

transient positive interest rates reflect subjects' willingness to borrow in order to execute the securities trades that are needed to optimize their portfolio. Because purchases are paid for in cash, and orders in one market cannot be made conditional on trades in other markets, subjects need to have a certain amount of cash to execute their purchase orders if they do not wish to sit there and wait for their sales orders for other securities to be executed first.

As a matter of fact, monetary theory has appealed to such a *cash in advance* constraints to explain the presence of interest-bearing nominally riskfree securities along with zero-interest cash (Clower 1967; Lucas 1982). Monetary theorists, however, have generally attributed cash-in-advance constraints to the purchase of consumption goods. In our experiments, they are clearly at work in investment decisions as well. Our findings would suggest that it may be worthwhile to study the effect of portfolio re-allocations in explaining the behavior of interest rates in naturally occurring financial markets as well.

Figure 2 displays the evolution of transaction prices for experiment 011126. Remember that Bulgarian subjects were used in this experiment,




Figure 2. Evolution of securities prices in the 01126 experiment (involving Bulgarian subjects). Vertical lines delineate periods; horizontal lines depict expected payoffs (top lines: security A; middle lines: security B; bottom lines: notes). All transactions of risky securities occurred at prices strictly below expected payoffs, suggesting risk aversion.

and that they were paid nominally as much as North American subjects, which means that they were paid effectively more than 10 times as much, because of the difference in standard of living. In this experiment, states were drawn without replacement. As a result, probabilities, and hence, expected payoffs on risky securities change each period. This is reflected in the discrete jumps in the horizontal lines. Again, prices are always below expected payoff; the risk premia are even bigger now than in the the first experiment, 981007. Presumably, this can be attributed to the more than tenfold increase in risk.¹⁸

For the last 10 transactions in each period and experiment, Table II lists the ratio of the average price over the expected payoff. (Only data for the risky securities are displayed.) One can verify that the visual evidence from Figures 1 and 2 is representative. With few exceptions, the ratios are all strictly below 1, confirming the presence of significant risk premia. A few

¹⁸ Holt and Laury (2002) also provide extensive experimental evidence of increases in risk premia as a result of increases in the size of risk.

Table II. Ratio of prices over expected payoffs (prices averaged over last 10 trades in any security; for securities that did not trade, previous transaction prices are taken)

Date	Security	Period									
		1	2	3	4	5	6	7	8	9	10
981007	A	0.96	0.93	0.93	0.95	0.95	0.88				
	B	0.96	0.98	0.97	0.97	0.97	0.97				
	Notes	0.95	0.97	0.99	0.97	0.99	0.99				
981116	A	0.93	0.88	0.91	0.92	0.82	0.88				
	B	0.94	0.97	0.98	0.97	0.95	0.94				
	Notes	0.99	1.00	0.98	0.99	1.00	0.99				
990211	A	0.96	0.99	0.93	0.88	0.95	0.97	1.00			
	B	0.92	0.91	0.93	0.89	0.94	0.91	0.97			
	Notes	0.94	0.95	0.96	0.98	0.96	0.99	0.97			
990407	A	0.95	0.90	0.90	0.87	0.87	0.92	0.88	0.90		
	B	0.98	0.99	1.01	1.03	1.08	1.01	1.02	1.08		
	Notes	0.99	0.99	1.00	0.99	0.99	0.99	0.99	0.99		
991110	A	0.89	0.92	0.93	0.93	0.91	0.90				
	B	0.84	0.86	0.90	0.95	0.96	0.95				
	Notes	0.96	0.97	0.97	0.98	0.98	1.00				
991111	A	0.98	0.94	0.98	0.98	1.00	1.01	0.92	0.90		
	B	0.98	1.00	0.91	0.91	0.94	0.95	0.95	0.95		
	Notes	0.99	0.99	0.99	0.99	0.98	0.99	0.99	0.99		
011114	A	0.98	0.92	0.89	0.89	0.99	0.98	0.99	1.00	0.91	0.93
	B	0.94	0.97	0.96	0.96	0.93	0.96	1.01	0.94	0.96	0.90
	Notes	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.98	0.99	0.97
011126	A	0.78	0.78	0.83	0.84	0.85	0.84	0.83	0.86		
	B	0.74	0.93	0.91	0.90	0.94	0.93	0.95	0.90		
	Notes	0.90	1.10	0.99	0.98	0.98	0.98	1.02	0.99		
011205	A	0.91	0.93	0.95	0.89	0.87	0.94	0.97	1.00	0.90	0.87
	B	0.97	0.93	0.92	0.93	0.90	0.94	0.94	0.95	0.96	0.92
	Notes	0.99	1.00	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99

exceptions occur in experiments 990407 and 991111. These exceptions can be attributed to mis-understanding (or disbelief) about the independence of the drawing of the states. We will return to this issue.¹⁹

¹⁹ Note that there is no evidence of a reduction in risk premia over time. As mentioned in Section 1, even if subjects are risk neutral, the option to maximize gains by remaining solvent makes subjects temporarily cautious. By the end of the experiment, however, risk neutrality kicks in and risk premia should disappear. Table II shows that risk premia are generally no lower in later periods.

4.2. EVIDENCE FOR THE STATE PRICE PROBABILITY RATIO PROPERTY

We deduced state prices for the three states from the transaction prices of the three securities. As has become standard in mathematical finance and derivatives analysis, we normalized the prices to add up to one. We then divide the state prices by the probabilities of the states to obtain the state price probability ratio. The State Price Probability Ratio Property predicts that state price probability ratios will be ranked inversely to the aggregate wealth in each state. In all periods of all experiments, initial endowments and payoffs were such that the aggregate wealth was highest in state Y and lowest in state X. Consequently, state price probability ratios should be lowest in Y and highest in X.

Figure 3 displays the evolution of the three state price probability ratios (one for each state) in experiment 981007. The state price probability ratios have the tendency to separate as predicted by the theory. The ratios eventually become highest for X and lowest for Y, exactly as required in the Arrow–Debreu equilibrium. Note that the state price probability ratios for X are generally above 1, indicating that state security X (implicitly) traded at a price *above* its expected payoff. That is, its risk premium is *negative*, a result that one may find puzzling at first. It is remarkable that markets do manage to price risky securities correctly even if the resulting prices are counter-intuitive.²⁰

As emphasized in Section 1, that the State Price Probability Ratio Property might appear in the data is far from obvious. Even if subjects had collectively wanted to price the state securities in accordance with the theoretical prediction, they would still have to know the aggregate endowment in each state. They were not given this information and the experimental set-up lacked sufficient transparency for them to deduce it from the trading data.

Figure 4 displays the evolution of the three state price probability ratios for the Bulgarian experiment, 011126. The separation is much earlier and pronounced than in 981007, presumably because of the significantly higher risk. The visual evidence is cleaner despite the fact that market prices, and hence, state prices, have to adjust to new probabilities each period.

In periods 1 and 2, state prices implied from transaction prices were sometimes negative, and hence, state price probability ratios ended up negative as well. In principle, this implies arbitrage opportunities. In practice, most of these arbitrage opportunities could only be exploited to a limited extent, for two reasons. First, transactions do not occur in synchrony, so some transaction prices may be obsolete. Second, short-sale constraints on the risky assets generally make it impossible to construct portfolios that

²⁰ In a pilot experiment where subjects were trading only state securities, subjects reported afterwards to have been puzzled when observing that one of the state prices moved above the expected payoff of the asset.




Figure 3. Evolution of ratios of state prices over state probabilities in the 981007 experiment. Vertical lines delineate periods. In the Arrow–Debreu equilibrium with expected utility preferences, the state price probability ratio of X should be highest; that of Y should be lowest.

mimic the payoff on the state securities, and hence, that exploit the negative prices.

One would like a formal test to confirm that state price probability ratios tend to move in the direction predicted by the Arrow–Debreu equilibrium. We propose here a test that determines whether state price probability ratios adjust correctly when their ranking is not as predicted by the theory. We take random walk pricing as our null hypothesis. The random walk hypothesis is an appropriate benchmark, because it is devoid of economic content beyond the statement that it does not allow for either arbitrage opportunities or simple speculative profit opportunities. Indeed, it is the best model of price behavior one had before asset pricing theory was developed. The tradition behind this model goes back to Bachelier (1900).

Our test calculates the probability of observing the dynamics in state price probability ratios actually observed in our experiments if price dynamics were merely of the random walk type. Under the alternative that markets are attracted by the Arrow–Debreu equilibrium, we expect specific changes in the




Figure 4. Evolution of ratios of state prices over state probabilities in the 011126 experiment (involving Bulgarian subjects). In the Arrow–Debreu equilibrium with expected utility preferences, the state price probability ratio of X should be highest; that of Y should be lowest.

state price probabilities when they are not aligned appropriately. In particular, we expect the following.

Ranking of state price probability ratios at t	Expected effect
$\rho_{X,t} > \rho_{Y,t} > \rho_{Z,t}$	$(\rho_{Z,t+1} - \rho_{Y,t+1}) - (\rho_{Z,t} - \rho_{Y,t}) > 0$
$\rho_{X,t} > \rho_{Z,t} > \rho_{Y,t}$	Anything is possible
$\rho_{Y,t} > \rho_{X,t} > \rho_{Z,t}$	$(\rho_{X,t+1} - \rho_{Y,t+1}) - (\rho_{X,t} - \rho_{Y,t}) > 0$
$\rho_{Y,t} > \rho_{Z,t} > \rho_{X,t}$	$(\rho_{X,t+1} - \rho_{Y,t+1}) - (\rho_{X,t} - \rho_{Y,t}) > 0$ or $(\rho_{Z,t+1} - \rho_{Y,t+1}) - (\rho_{Z,t} - \rho_{Y,t}) > 0$
$\rho_{Z,t} > \rho_{Y,t} > \rho_{X,t}$	$(\rho_{X,t+1} - \rho_{Y,t+1}) - (\rho_{X,t} - \rho_{Y,t}) > 0$
$\rho_{Z,t} > \rho_{X,t} > \rho_{Y,t}$	$(\rho_{X,t+1} - \rho_{Z,t+1}) - (\rho_{X,t} - \rho_{Z,t}) > 0$

These predictions are weak, because they only concern the *sign* of the change in the difference between two state price probabilities. The question

is: are economic forces strong enough that the expected effects can be detected sharply?

As test statistic, we compute the frequency of observing (transisting to) the expected outcome for each state (ranking of state price probabilities). We subsequently take the mean across states. The latter is mandated by the fact that, in finite samples, not all states need occur, in which case some transition frequencies are undefined. Notice that the second frequency will always be 1.²¹ We include this frequency, so that outcomes where the Arrow–Debreu prediction holds ($\rho_{X,t} > \rho_{Z,t} > \rho_{Y,t}$) receive more weight. Let τ denote this *mean transition frequency*.

To test whether the observed mean transition frequencies are unusual under the assumption of a random walk, and hence, to determine whether to reject the null of a random walk, we compute the distribution of τ under the null hypothesis by bootstrapping the empirical joint distribution of price changes in each experiment. In the bootstrap, we generated 200 price series of the same length as the sample used to estimate τ .²² We reject the null if the observed mean transition frequency is in the tails of the bootstrapped distribution for τ .

Table III reports the results. In addition to the observed mean transition frequencies (second column), we report 5%, 90% and 95% bootstrapped critical values under the null of random walk pricing. The striking aspect of the estimated mean transition frequencies is not only that they are highly significant, but also that they are remarkably constant across experiments. There are pronounced differences across experiments in terms of beginning prices and empirical distribution of price changes, which translate into marked differences in the distribution of τ under the null of a random walk.²³ Across all experiments, the estimated values for τ are almost always the same, however, suggesting that the same forces are at work. They are in the right tail of the distribution under the null, and, except for 991110, significant at the 5% level. Overall, Table III provides formal evidence that Arrow–Debreu equilibrium predicts price movements in experimental financial markets better than the random walk hypothesis.

²¹ The expected outcome is “anything is possible.” Consequently, we always declare to have expected the actual outcome when the second state occurs.

²² We bootstrapped the *mean-corrected* empirical distribution, in order to stay with the null hypothesis of a random walk.

²³ If the initial price configuration satisfies or is close to satisfying the Arrow–Debreu equilibrium restriction, the simulated τ s will be high (in the table above, the predicted outcome if $\rho_{X,t} > \rho_{Z,t} > \rho_{Y,t}$ obtains with unique frequency). This explains the high level of the 5% critical value for some experiments (990407 and 991110). Our procedure therefore penalizes experiments that happen to start out with prices that (closely or exactly) satisfy Arrow–Debreu equilibrium restrictions.

Table III. Test of random walk price evolution against Arrow–Debreu pricing theory

Experiment	Mean transition probability τ^a			
	Estimate	Critical value ^b		
		5%	90%	95%
981007	0.88**	0.31	0.59	0.61
981116	0.87**	0.25	0.48	0.51
990211	0.84**	0.39	0.79	0.83
990407	0.89**	0.64	0.83	0.85
991110	0.88*	0.64	0.88	0.89
991111	0.87**	0.36	0.70	0.73
011114	0.84**	0.28	0.52	0.54
011126	0.90**	0.21	0.77	0.78
011205	0.85**	0.54	0.79	0.85

^aFrequency that rank ordering of state price probability ratios is changed in the direction of the Arrow–Debreu equilibrium.

^bBased on 200 bootstrapped samples of the same size as used to estimate τ .

* = Significant at the 90% level. Significance levels are for the null of a random walk against the alternative of evolution towards the Arrow–Debreu equilibrium.

** = Significant at the 95% level.

It may deserve emphasis that the rejections of the random walk hypothesis reported in Table III do not imply that the subjects ignored profit opportunities from speculating on price changes. This is because our rejection of the random walk is based on information that subjects could not condition on, namely, the required rank ordering of state price probability ratios under the Arrow–Debreu equilibrium. The equilibrium ranking ratio is not readily determined from the history of price movements. Computation requires knowledge of the distribution of aggregate wealth across states. As emphasized before, subjects did not know this.

4.3. EVIDENCE FOR THE SHARPE RATIO PROPERTY

For experiments 981107 and 011126, respectively, Figures 5 and 6 display the evolution of the difference between the Sharpe ratio (mean excess return divided by standard deviation) of the market portfolio against the maximum possible Sharpe ratio.²⁴ Each observation corresponds to a trade in one of the three securities (the price of the other two securities was set equal to their last

²⁴ In computing the maximum Sharpe ratio, short-sale constraints on Securities A and B were accounted for.




Figure 5. Evolution of the distance of prices from CAPM equilibrium in the 981007 experiment. Distance is measured as the difference between the Sharpe ratio of the market portfolio and the maximal Sharpe ratio.

trading price). In the CAPM equilibrium, the distance between the market and maximal Sharpe ratio is zero. There is obviously a pronounced tendency for the markets to move toward the CAPM equilibrium, both across entire experiments and within a period (remember that in 011126, probabilities change, so that markets have to find a new CAPM equilibrium each period). Consequently, the experimental markets tend towards the Sharpe Ratio Property. As with the State Price Probability Ratio Property, we would like to emphasize again that there is an important sense in which this support for the CAPM is remarkable. Subjects are never told what the market portfolio is, so they cannot use the CAPM to price securities. Yet, CAPM pricing obtains. In other words, the emergence of the CAPM should be attributed to financial market equilibration in the presence of risk averse investors, and not to subjects' deliberate attempts to generate the results that support prevailing theory (the CAPM).

One can translate the visual evidence into formal statistical statements. In analogy with our test of the State Price Probability Ratio Property, we will test the null of a random walk against the alternative of equilibration towards CAPM. Note that even if prices form a random walk, markets can acci-




Figure 6. Evolution of the distance of prices from CAPM equilibrium in the 011126 experiment (involving Bulgarian subjects). Distance is measured as the difference between the Sharpe ratio of the market portfolio and the maximal Sharpe ratio.

mentally move towards the CAPM equilibrium, i.e., to a point where the market portfolio is mean–variance optimal. Under the alternative hypothesis, however, this movement is not accidental: the farther away from equilibrium, the more the market is attracted towards equilibrium. We wish to use this insight to formally distinguish the two hypotheses, and to test whether the visual evidence in favor of the CAPM is not a case of mere luck.

Our test works as follows. Let Δ_t denote the distance between the Sharpe ratio of the market and the maximum Sharpe ratio at transaction t . Δ_t is the absolute value of the variable plotted in Figures 5 and 6. Consider the projection of the change in Δ_t onto Δ_{t-1} :

$$\Delta_t - \Delta_{t-1} = \kappa \Delta_{t-1} + \epsilon_t, \quad (6)$$

where κ is such that ϵ_t is uncorrelated with Δ_{t-1} . CAPM implies $\Delta_t = 0$; convergence to CAPM pricing implies $\kappa < 0$. We then determine the distribution of the least squares estimates of κ under the null hypothesis of a random walk, by bootstrapping the empirical joint distribution of changes in

transaction prices. The null hypothesis of a random walk is rejected (in favor of stochastic convergence to CAPM) if the least squares estimate of κ is beyond a critical value in the left tail of the ensuing distribution. This testing procedure is a variation of *indirect inference* (see Gouriéroux et al., 1993): we summarize the data in terms of a simple statistical model (in our case, a least squares projection) and determine the distribution of the estimates by simulating the variables entering the statistical model. Instead of simulating off a theoretical distribution, we bootstrap the empirical distribution, however.²⁵

For each experiment, we estimated κ using OLS. 5% and 10% critical values under the random walk null hypothesis were determined by bootstrapping from the empirical joint distribution of price changes (we generated 200 price series of the same length as the sample used to estimate κ).²⁶

Table IV reports the results. The null of a random walk is rejected in each experiment. There is one exception: 991110. The latter may seem surprising when one inspects the graphical evidence, displayed in Figure 7. Visually, there seems to be a strong case for the CAPM. Yet, the statistical analysis indicates that there is a high probability to accidentally obtain a plot like the one in Figure 7, even if prices were a random walk. The contrast between the result of the statistical test and the visual evidence underscores the need for an appropriate, formal test of the theory, even in an experimental setting.

Rejection of the random walk hypothesis does not imply that subjects ignored profit opportunities. We reject the random walk on the basis of information that subjects did not have: the position of the market portfolio in mean–variance space.

4.4. CONTROL OF BELIEFS

Proper interpretation of the experimental results require that we have good control over subjects' beliefs as well as an informed guess about their attitudes towards risk. The latter is easiest, because the Arrow–Debreu equilibrium does not require that one know risk aversion; it can even change over time. The CAPM equilibrium requires that preferences can be approximated by quadratic utility, but, given the size of risk in our experiments, this is unlikely to be a major issue. Control of beliefs, however, is of major concern.

In particular, in two instances during the first set of experiments, prices move away from equilibrium in a direction that reflects misunderstanding (or disbelief) of the independence of the drawing process. The first occurrence is

²⁵ Gallant and Tauchen (1996) also use indirect inference, but, instead of matching a convenient statistical model, they match the scores of the likelihood function.

²⁶ We again bootstrapped the mean-corrected empirical distribution, in order to stay with the null hypothesis of a random walk.

Table IV. Test of random walk price evolution against CAPM pricing

Experiment	Attraction coefficient κ^a		
	Estimate	Critical value ^b	
		5%	10%
981007	-0.103**	-0.013	-0.009
981116	-0.081**	-0.009	-0.004
990211	-0.462**	-0.219	-0.081
990407	-0.016**	-0.009	-0.006
991110	-0.013	-0.119	-0.033
991111	-0.056**	-0.014	-0.007
011114	-0.042**	-0.007	-0.005
011126	-0.046**	-0.010	-0.001
011205	-0.037**	-0.007	-0.005

^aSlope coefficient in projection of changes in the difference between market Sharpe ratio and maximal Sharpe ratio onto past level.

^bBased on 200 bootstrapped samples of the same size as used to estimate κ .

** = Significant at the 5% level.

Significance levels are for the null of a random walk against the alternative of evolution towards the CAPM equilibrium.

in experiment 990407. The plot of the evolution of state prices (normalized to add up to 1) will make this clear. Figure 8 shows that the price of state Y is generally lowest, but that the ranking of states X and Z is wrong during most of the middle part of the experiment. An explanation for the reversal of the ordering of the prices for states X and Z emerges once it is contrasted with the sequencing of the states. The sequence of states in 990407 is: X (period 1), Y, Y, X, Z, Y, X, Y (period 8). By the fourth period, some subjects evidently become convinced that period Z is more likely to be drawn next, because it has not yet appeared (of course, this is incorrect). These subjects implicitly bid up the price of state Z. Their view of the world (that the likelihood of a state depends inversely on the number of times it is drawn in the past) seems to be confirmed at the end of period 5. As a result, by period 8, speculation about Z re-emerges. Again, they implicitly bid up the price of state Z, making it hard for the market to move to an Arrow–Debreu equilibrium.

The second instance in which subjective beliefs deviated from objective probabilities is around the middle of experiment 991110. The evolution of state prices is plotted in Figure 9. Notice how the price of state X is generally highest, as predicted in the Arrow–Debreu equilibrium, but that the ranking of the prices of states Y and Z is wrong in periods 3–6. Again this may be




Figure 7. Evolution of the distance of prices from CAPM equilibrium in the 991110 experiment. Distance is measured as the difference between the Sharpe ratio of the market portfolio and the maximal Sharpe ratio.

explained in terms of speculation. The sequence of states in the 991111 experiment is: Y (period 1), Z, Z, X, X, Y, Z, Z (period 8). From period 4 on, subjects speculate that Z is less likely, whereas Y, and to a lesser extent, X, are more likely. Consequently, the prices of Y and X both increase.

In both instances, some subjects evidently misunderstand (or distrust) independence. Instead, they hold beliefs whereby the likelihood of a state changes inversely with the number of past occurrences. That is, they seem to hold hypergeometric beliefs, whereby states are drawn *without replacement*.²⁷ Consequently, in order to have better control over subjects' beliefs, we switched procedures in the second set of experiments (011114, 011126 and 011205). States were drawn without replacement, from an initial distribution of four balls of type X, four of type Y, and four of type Z. No more anomalies emerged. As a matter of fact, convergence to equilibrium was no

²⁷ When computing state price probability ratios and Sharpe ratios based on hypergeometric beliefs, the anomalies in experiments 990407 and 991111 disappeared. Details can be obtained from the authors upon request.




Figure 8. Evolution of ratios of state prices over state probabilities in the 990407 experiment. Vertical lines delineate periods. In the Arrow–Debreu equilibrium with expected utility preferences, the state price probability ratio of X should be highest; that of Y should be lowest. Rank ordering is violated during much of the middle of the experiment, presumably because of mis-understanding (or disbelief) that state probabilities did not change over time.

worse, despite the increase in complexity (markets had to search for a new equilibrium each period if only because probabilities changed).²⁸

A final comment about control of beliefs is warranted. The anomalies in experiments 990407 and 991111 hardly affected the formal evidence against the random walk and in favor of the Arrow–Debreu equilibrium (Table III) or the CAPM (Table IV). At the margin at least (i.e., at the level of transaction price changes), there remains a tendency for prices to adjust in the direction of Arrow–Debreu or CAPM equilibrium even in the presence of speculative bounds.

²⁸ It is not clear how many subjects mis-understood independence. It may be that some did not believe that nobody had privileged information about the state to be drawn, despite claims to the contrary in the instructions. Such subjects would attempt to read this privileged information from prices, thus amplifying anomalies.




Figure 9. Evolution of ratios of state prices over state probabilities in the 991111 experiment. Vertical lines delineate periods. In the Arrow–Debreu equilibrium with expected utility preferences, the state price probability ratio of X should be highest; that of Y should be lowest. Rank ordering is violated during much of the middle of the experiment, presumably because of mis-understanding (or disbelief) that state probabilities did not change over time.

4.5. EFFECT OF THE SOLVENCY RULE

As mentioned in Section 1, our solvency rule induces subtle dynamics in optimal choices, and therefore, could potentially have affected prices. Table II demonstrates, however, that there are no significant trends in risk premia (discounts relative to expected payoffs). Furthermore, it is possible to demonstrate that even individual choices hardly show any sign of dynamics, as follows.

Take the first set of experiments. Periods are independent replications of each other, because states are drawn with replacement. From the perspective of subjects, only three facets change across periods: (i) relative prices, (ii) personal wealth and (iii) time. The latter two could affect choices if they lead to systematic changes in preferences, most importantly, through our solvency rule. Otherwise, effectively only one element changes from period to period, namely relative prices. If changes in personal wealth or time do not induce preference dynamics, choices must satisfy the axioms of revealed preference

(provided that they are optimal, of course). We now show that actual choices indeed never violate the axioms of revealed preferences for 1/4 of the subjects, and that they violate these axioms insignificantly in most other cases. In other words, in the vast majority, subjects' choices can be almost perfectly rationalized as if coming from a single (quasiconcave) utility function unaffected by changes in wealth and/or time. This also means that our solvency rule does not induce significant dynamics in choices.²⁹

Our results are based on an extension of methodology suggested by Andreoni and Miller (2002). We use ideas originally introduced by Afriat (1973). Let $t = 1, \dots, T$ index periods. Let p_t denote the vector of prices of securities at the end of period t (the price of the riskfree security is set equal to 100). Let x_t denote the vector of a subject's securities holdings at the end of period t (cash holdings are converted to holdings of riskfree securities at the price of 100 francs per unit). A subject's choices can be rationalized as the constrained maximization of some quasiconcave utility function if there exist numbers $U_t \geq 0$ and $\lambda_t \geq 1$ such that³⁰

$$U_t - U_{t'} \leq \lambda_{t'} p_{t'}(x_t - x_{t'}), \quad (7)$$

for $t, t' = 1, \dots, T, t \neq t'$.

We propose to measure *failure* by means of a slack variable ε . If $\varepsilon = 0$, the Afriat inequalities hold. If $\varepsilon > 0$, they fail, and ε is the minimum income that has to be taken away from the subject in at least one period (but at most $T - 1$ periods) for the Afriat inequalities to hold. ε solves the following non-linear programming problem.

$$\min \varepsilon, \quad (8)$$

²⁹ It may be surprising that preferences appear to be fixed during the course of an experiment despite our solvency rule. We are not claiming that subjects ignored the solvency rule. Instead, we are showing evidence that the effect of the solvency rule is insignificant. An analogy with experiments in physics is appropriate. Newtonian physics makes predictions about the movement of objects in a vacuum. On earth, it is impossible to create a perfect vacuum, because friction is always present. In the laboratory, however, it is possible to create conditions such that Newtonian physics makes valid predictions despite the presence of friction. We generated an analogous situation in our financial markets experiments. Static asset pricing theory explains the experimental results despite the fact that it ignores the preference dynamics that our solvency rule induces.

³⁰ In Afriat's original formulation, the restriction on λ_t is strict positivity ($\lambda_t > 0$). Inspection of the Afriat inequalities, however, reveals that the scale of the U_t s and λ_t s is arbitrary, so that we can require $\lambda_t \geq 1$. At the same time, this ensures that λ_t is strictly positive. The constraint on λ_t is thereby converted to a weak inequality, which facilitates numerical verification of the inequalities.

$$\begin{aligned}
U_{t'} - U_t &\geq \lambda_{t'} [p_{t'} \cdot x_{t'} - \varepsilon - p_t \cdot x_t] \quad t, t' = 1, \dots, T; t \neq t', \\
\varepsilon &\geq 0, \\
\lambda_{t'} &\geq 1 \quad t' = 1, \dots, T, \\
U_t &\geq 0 \quad t = 1, \dots, T.
\end{aligned}$$

Table V provides details of the distribution of violations of the Afriat inequalities for the six experiments in the first set. Of the 243 total of active subjects (two subjects never changed positions and were excluded from the analysis), one quarter (62) do *not* violate any of the Afriat inequalities. The majority (62 + 79) have violations of no more than 25 US cents. This means that choices of the majority of subjects can be rationalized in terms of a time-invariant quasiconcave preference profile provided one is willing to shift the value of their holdings in at least one period by at most 25 US cents. This is insignificant relative to the typical total value of their holdings in any period, which is more than US \$50. Very few subjects (13) have violations of more than US \$3. The highest violation is US \$5.76.³¹

In evaluating violations of revealed preference, one has to keep in mind that they have been computed assuming that subjects could trade at last transaction prices while subjects' actual choices may have been based on slightly different prices. In addition, the anomalous belief dynamics reflected in prices in some of the experiments (see Section 4.4) must have their origin in anomalous choices at the individual level, and hence, could also explain some of the violations. Finally, the above exercise is limited to the first set of experiments. In the second set of experiments, not only do relative prices change, but also probabilities. Choices therefore change not only because prices change, but also because beliefs change. Hence, axioms of revealed preference can no longer be checked.

³¹ To get a further feeling for the meaning of violations of the Afriat inequalities, take experiment 990407, for instance. Consider a subject who is endowed with the market portfolio (expressed per capita) and 400 francs cash. Take the relative prices at the end of periods 1 and 5 (see Table 2). Imagine that this subject invests all wealth in security A when it was most expensive (period 1), while rebalancing all wealth into security B when it became most expensive (at the end of period 5). To do this, this subject sells all other securities, and borrows money (sells short Notes) up to the short-sale constraint (8 Notes). Such an investment policy may be chosen if the subject erroneously thinks that the higher the price of an asset, the more likely it will pay its highest possible dividend. Clearly, this subject violates the Afriat inequalities. The slack variable that solves the non-linear programming problem in Equation (8) equals 282 francs, or about US \$10. No subject in our experiments generated a slack variable as high as this.

Table V. Frequency of violations of afriat inequalities – experiments with state drawing of type I

Size of Violations (ε), in U.S. dollar ^a					
$\varepsilon = 0$	$0 < \varepsilon \leq \text{US } \0.25	$\text{US } \$0.25 < \varepsilon \leq \text{US } \1	$\text{US } \$1 < \varepsilon \leq \2	$\$2 < \varepsilon \leq \3	$\$3 < \varepsilon$
62 ^b	79	44	31	14	13

^a ε in the non-linear programming problem in (8).

^bNumber of subjects in the category $\varepsilon = 0$.

4.6. EFFECT OF MARKET THICKNESS

Relative to standards in experimental economics, the experiments we have been discussing here are large-scale. Previously, our own (CAPM) experiments involved only 5–13 subjects; see Bossaerts and Plott (2002). While these numbers are generally considered to be sufficient to generate Walrasian equilibrium outcomes, such consensus is based on much simpler settings: one good/security and no uncertainty. In contrast, here we have three securities, and we let subjects' natural attitudes towards uncertainty determine pricing. This appears to have a marked effect on pricing. The results of the thin-markets experiments in Bossaerts and Plott (2002) are substantially less compelling as far as general equilibrium theory is concerned.

To illustrate how different the results are, Table VI displays medians³² of the distance from CAPM pricing for: (i) the seven thin-market experiments in Bossaerts and Plott (2002) (where there were 5–13 subjects), (ii) six thick-market experiments reported on here (where there were between 19 and 63 subjects).³³ The distance from CAPM pricing is measured as the difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio. Standard deviations (of the medians) across experiments are in parentheses. With the exception of the first two periods, the thick-market experiments produce pricing that is significantly closer to CAPM pricing.

5. Conclusion

The experiments demonstrate that the basic principles of modern asset pricing theory are very much at work in competitive financial markets. Prices tend towards a pattern where risk premia are determined solely by covariance

³² If the median fell between two outcomes, the *lower* of the two outcomes is reported.

³³ Results for the six experiments in the first set are shown here, because only those are directly comparable to the ones in Bossaerts and Plott (2002), where states were always drawn with replacement.

Table VI. Distance from CAPM equilibrium: comparison with thin-markets experiments

	Period							
	1	2	3	4	5	6	7	8
Thin markets ^a	0.17 ^b (0.15)	0.10 (0.05)	0.18 (0.10)	0.17 (0.10)	0.14 (0.10)	0.15 (0.03)	0.12 (0.16)	0.14 (0.06)
Thick markets ^c	0.20 (0.18)	0.07 (0.07)	0.06 (0.06)	0.07 (0.07)	0.04 (0.02)	0.08 (0.30)	0.06 (0.09)	0.03 (0.02)

^aBased on the seven experiments reported in Bossaerts and Plott (2002).

^bMedian (or one smaller if the median is not equal to a realized outcome) of the average difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio. Corresponding standard deviations in parentheses.

^cBased on the six experiments of the first set reported in this paper.

with aggregate risk. These results are rendered more significant because subjects lacked the information that was necessary to deliberately set prices according to this prediction.

A judicious choice of trading interface and number of subjects enabled us to emulate competitive markets in the laboratory. The success of the Arrow–Debreu model and the CAPM in explaining prices suggests that our conjecture about preferences was correct (smooth expected utility and quadratic utility, respectively). That the CAPM emerges is probably attributable to the relatively small levels of risk, which allows one to approximate utility with an expansion that only includes linear and quadratic terms.

In contrast, beliefs are more difficult to control. We conjectured that a number of subjects either mis-understand or disbelieve statistical independence. The second set of experiments, where states were drawn without replacement, confirmed the conjecture, as transient anomalous pricing behavior disappeared.

Our study has focused on pricing. Asset pricing theory also makes predictions about allocations. These can be studied using the experiments reported on in this paper. An example of the type of information one can extract from holdings was in Section 4.5, where we studied whether subjects' choices could be rationalized in terms of a time-invariant quasiconcave preference profile. Additional theoretical predictions can be investigated, but since such studies raise new issues they are best discussed in a separate paper. One of the most important questions that comes to mind is whether prices are determined by the choices (holdings) of an identifiable marginal investor, as is often stated in the finance literature, or whether all subjects are marginal, each contributing to pricing, as predicted by general equilibrium theory. Bossaerts, Plott and Zame (2003) examine this question.

Finally, there is an urgent need to understand how prices manage to move to equilibrium. The present paper demonstrates that prices move in the right direction, leaving open the question of why. In particular, what role do agents' choices play? A first exploration can be found in Asparouhova, Bossaerts and Plott (2003). An explicit model and its implications for empirical analysis of asset pricing theory on field data can be found in Bossaerts (2003). Elements of the theory proposed in Crockett, Spear and Sunder (2002) may provide insights of how repetition of identically the same environment (as in our first set of experiments) eventually leads to equilibrium pricing.

References

- Afriat, S. N. (1973) On a system of inequalities in demand analysis: An extension of the classical method, *International Economic Review* **14**, 460–472.
- Andreoni, J. and Miller J. (2002) Giving according to GARP: An experimental test of the consistency of preferences for Altruism, *Econometrica* **70**, 737–753.
- Asparouhova, E., Bossaerts, P., and Plott, C. (2003) Excess demand and equilibration in multi-security financial markets: The empirical evidence, *Journal of Financial Markets* **6**, 1–22.
- Bachelier, L. (1900) Théorie de la Spéculation, *Annales Scientifiques de l' Ecole Normale Supérieure* III-17, 21–86.
- Barberis, N., Huang, M., and Santos, T. (2001) Prospect theory and asset prices, *Quarterly Journal of Economics* **116**, 1–53.
- Biais, B., Bossaerts, P., and Spatt, C. (2003) Equilibrium asset pricing under heterogeneous information, CMU Working Paper.
- Bossaerts, P. (2003) Testing CAPM in real markets: Implications from experiments, Caltech Working Paper.
- Bossaerts, P., Fine, L., and Ledyard, J. (2002) Inducing liquidity in thin financial markets through combined-value trading mechanisms, *European Economic Review* **46**, 1671–1695.
- Bossaerts, P., Kleiman, D., and Plott, C. (2002) Price discovery in financial markets: The case of the CAPM, in (ed.), C. Plott *Collected papers on the Foundations of Experimental Economics and Political Science: Information, Finance and General Equilibrium*, Volume. III, Edwin Elgar Publishers.
- Bossaerts, P., Plott, C., and Zame, W. (2003), Prices and portfolio choices in financial markets: theory and experiments, Caltech/UCLA Working Paper.
- Bossaerts, P. and Plott, C. (2002) The CAPM in thin experimental financial markets, *Journal of Economics Dynamics and Control* **26**, 1093–1112.
- Clower, R. (1967) A re-consideration of the microfoundations of monetary theory, *Western Economic Journal* **6**, 1–9.
- Crockett, S. Spear, S., and Sunder, S. (2002) A simple decentralized institution for learning competitive equilibrium, Yale ICF Working Paper.
- Debreu, G. (1959): *Theory of Value*, New York, Wiley.
- Dybvig, P. and Ingersoll, J. (1982) Mean–variance theory in complete markets, *Journal of Business* **55**, 233–251.

- Epstein, L. and Zin, S. (2001) The independence axiom and asset returns, *Journal of Empirical Finance* **8**, 537–572.
- Gallant, R.A. and Tauchen, G. (1996) Which moments to match? *Econometric Theory* **12**, 657–681.
- Gouriéroux, C., Monfort, A., and Renault, E. (1993) Indirect inference, *Journal of Applied Econometrics* **8**, S85–S118.
- Holt, C. and Laury, S. (2002) Risk aversion and incentive effects, *American Economic Review* **92**, 1644–1655.
- Kroll, Y., Levy, H., and Rapaport, A. (1988) Experimental tests of the separation theorem and the capital asset pricing model,” *American Economic Review* **78**, 500–519.
- Levy, H. (1997) Risk and return: An experimental analysis, *International Economic Review* **38**, 119–149.
- Lintner, J. (1965) The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* **47**, 13–37.
- Lucas, R. (1982) Interest rates and currency prices in a two-country world, *Journal of Monetary Economics* **10**, 335–59.
- Mossin, J. (1966) Equilibrium in a capital asset market *Econometrica* **34**, 768–83.
- Noussair, C., Plott, C. and Riezman, R. (1997) The principles of exchange rate determination in an international Finance experiment, *Journal of Political Economy* **105**, 822–861.
- Roll, R. (1977) A critique of the asset pricing theory’s tests, Part I: On the past and potential testability of the theory, *Journal of Financial Economics* **4**, 129–76.
- Sharpe, W. (1964) Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* **19**, 425–442.

Copyright of Review of Finance is the property of Kluwer Academic Publishing / Business and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.