Risk Aversion in Laboratory Asset Markets

Peter Bossaerts*

California Institute of Technology
Centre for Economic Policy Research

William R. Zame†

UCLA
California Institute of Technology

March 15, 2005

*Financial support from the R. G. Jenkins Family Fund and the National Science Foundation is gratefully acknowledged. Opinions, findings, conclusions and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agency.

†Financial support from the John Simon Guggenheim Memorial Foundation, the National Science Foundation, the Social and Information Sciences Laboratory at Caltech, and the UCLA Academic Senate Committee on Research is gratefully acknowledged. Opinions, findings, conclusions and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agency.

Abstract

This paper reports findings from a series of laboratory asset markets. Although stakes in the experiment are modest, the data display clear evidence of substantial risk aversion. Most obviously, asset prices imply a substantial equity premium: risky assets are priced substantially below their expected payoffs. Moreover, the differences between expected asset payoffs and asset prices are in the direction predicted by standard asset-pricing theory: assets with higher beta have higher returns. The data yield estimatse of the Sharpe ratio of the market in the range $0.2 - 1.7$ (the Sharpe ratio of the New York Stock Exchange is approximately .43), and CAPM yields estimates of the market absolute risk aversion on the order of 10^{-3}. This work suggests useful ways to separate the effects of risk aversion from competing explanations in other experimental environments.

1 Introduction

Forty years of econometric tests have provided only weak support for the predictions of asset pricing models. (See Davis, Fama & French (2000) for instance.) However, it is difficult to know where the problems in such models lie, or how to improve them, because basic parameters of the theories — including the market portfolio, the true distribution of asset returns, the information available to investors — cannot be observed in the historical record. Laboratory tests of these theories are appealing because these basic parameters (and others) can be observed accurately — or even controlled. However, most asset pricing theories rest on the assumption that individuals are risk averse.[1] Because risks and rewards in laboratory experiments are (almost of necessity) small (in comparison to subjects' lifetime wealth, or even current wealth), the degree of risk aversion observable in the laboratory might be so small as to be undetectable in the unavoidable noise, which would present an insurmountable problem.

This paper reports findings from a series of laboratory asset markets that bely this concern: despite relatively small risks and rewards, the effects of risk aversion are detectable and significant. Most obviously, asset prices imply a significant equity premium: risky assets are priced significant below their expected payoffs. Moreover, the differences between expected asset payoffs and returns (payoffs per unit of investment) are in the direction predicted by standard asset-pricing theory: assets with higher beta have higher returns. As a quantitative expression of the degree of risk aversion, we obtain estimates of Sharpe ratios of the market in the range $0.2-1.7$ (the Sharpe ratio of the NYSE is approximately 0.43), and, using CAPM, we estimate the market absolute risk aversion to be approximately 10^{-3}. Our work suggests useful ways to distinguish the effects of risk aversion from subject errors, quantal response equilibrium, etc. in a number of experimental environments.

[1]Here we refer to theories such as the Capital Asset Pricing Model of Sharpe (1964) that predict the prices of fundamental assets, rather than to theories such as the pricing formula of Black & Scholes (1973) that predict the prices of options or other derivative assets. The latter theories do not rest on assumptions about investor risk attitudes, but rather on the absence of arbitrage.

In our laboratory markets, 30 - 60 subjects trade one riskless and two risky securities (whose dividends depend on the state of nature) and cash. Each experiment is divided into 6-9 periods. At the beginning of each period, subjects are endowed with a portfolio of securities and cash. During the period, subjects trade through a continuous, web-based open-book system (a form of double auction that keeps track of infra-marginal bids and offers). After a pre-specified time, trading halts, the state of nature is drawn, and subjects are paid according to their terminal holdings. The entire situation is repeated in each period but the state of nature is drawn anew at the end of each period. Subjects know the dividend structure (the payoff of each security in each state of nature) and the probability that each state will occur, and of course they know their own holdings and their own attitudes toward wealth and risk. They also have access to the history of orders and trades. Subjects do not know the number of participants in any given experiment, nor the holdings of other participants, nor the market portfolio.

Typical earnings in a single experiment (lasting 2+ hours) are $50-100 per subject. Although this is a substantial wage for some subjects, it is small in comparison to lifetime wealth, or indeed to current wealth (the pool of subjects consists of undergraduates and MBA students). Small rewards suggest approximately risk neutral behavior, asset prices nearly coincident with expected payoffs, little incentive to trade, and hence little trade at all.

However, our experimental data are inconsistent with these implications of risk neutrality; rather the data suggest significant risk aversion. Most obviously, market prices are below expected returns, and substantial trade takes place. Moreover, assets with higher beta have higher returns (lower prices), as suggested by standard asset pricing theories. Quantitative measures of risk aversion are provided by the Sharpe ratios of the market portfolio, which are in the range $0.2 - 1.7$ — on the same order as the Sharpe ratio of the New York Stock Exchange (computed on the basis of yearly data), which is 0.43 — and the imputed market risk aversion derived from CAPM, which is approximately 10^{-3}.

Following this Introduction, Section 2 describes our experimental asset markets, Section 3 presents the data generated by these experiments and the relationship of these data to standard asset pricing theories. Section 4 suggests implications of our experiments for the design and interpretation of other experiments where risk aversion may play a role, and concludes.

2 Experimental Design

In our laboratory markets the objects of trade are *assets* (state-dependent claims to wealth at the terminal time) A, B, N (*Notes*) and *Cash*. Notes are riskless and can be held in positive or negative amounts (can be sold short); assets A, B are risky and can only be held in non-negative amounts (cannot be sold short).

Each experimental session of approximately 2 hours is divided into 6-9 *periods*, lasting 15-20 minutes. At the beginning of a period, each subject (investor) is endowed with a portfolio of assets and Cash; the endowment of risky assets and Cash are non-negative, the endowment of Notes is negative (representing a loan that must be repaid). During the period, the market is open and assets may be traded for Cash. Trades are executed through an electronic open book system (a continuous double auction). During the period, while the market is open, no information about the state of nature is revealed, and no credits are made to subject accounts; in effect, consumption takes place only at the close of the market. At the end of each period, the market closes, the state of nature is drawn, payments on assets are made, and dividends are credited to subject accounts. (In some experiments, subjects were also given a bonus upon completion of the experiment.) Accounting in these experiments is in a fictitious currency called *francs*, to be exchanged for dollars at the end of the experiment at a pre-announced exchange rate. Subjects whose cumulative earnings at the end of a period are not sufficient to repay their loan are bankrupt; subjects who are bankrupt for two consecutive trading periods are barred from trading in future periods.[2] In effect, therefore, consumption in a given period can be negative.

Subjects know their own endowments, and are informed about asset payoffs in each of the 3 states of nature X, Y, Z, and of the objective probability distribution over states of nature. We use two treatments of uncertainty. In the first treatment, states of nature for each period are drawn independently with probabilities $1/3, 1/3, 1/3$; randomization is achieved by using a random number generator or by drawing with replacement from an urn containing

[2] However, the bankruptcy rule was seldom triggered.

equal numbers of balls representing each state. In the second treatment, balls, marked with the state, are drawn *without replacement* from an urn initially containing 18 balls, 6 for each state. (Subjects are informed of the procedure.) Asset payoffs are shown in Table 1 (1 unit of Cash is 1 franc in each state of nature), and the remaining parameters for each experiment are shown in Table 2. (Experiments are identified by year-month-day.)

In all experiments, subjects were given complete instructions, including descriptions of some portfolio strategies (but no suggestions as to which strategies to choose). Complete instructions and other details are available at http//eeps3.caltech.edu/market-011126; use anonymous login, ID 1, password a.

Table 1: Asset Payoffs

State	X	Y	Z
A	170	370	150
B	160	190	250
N	100	100	100

Subjects are *not* informed of the endowments of others, or of the market portfolio (the social endowment of all assets), or the number of subjects, or whether these are the same from one period to the next. The information provided to subjects parallels the information available to participants in stock markets such as the New York Stock Exchange and the Paris Bourse. We are especially careful not to provide information about the market portfolio, so that subjects cannot easily deduce the nature of aggregate risk — lest they attempt to use a standard model (such as CAPM) to *predict* prices, rather than to take observed prices as given. Keep in mind that neither general equilibrium theory nor asset pricing theory require that participants have any more information than is provided in these experiments. Indeed, much of the power of these theories comes precisely from the fact that agents know *only* market prices and their own preferences and endowments.

Keep in mind that the social endowment (the market portfolio), the distribution of endowments, and the set of subjects and hence preferences differ

Table 2: Experimental Parameters

Date	Draw Type [a]	Subject Category (Number)	Bonus Reward (franc)	Endowments A	Endowments B	Endowments Notes[b]	Cash (franc)	Exchange Rate $/franc
981007	I	30	0	4	4	-19	400	0.03
981116	I	23	0	5	4	-20	400	0.03
		21	0	2	7	-20	400	0.03
990211	I	8	0	5	4	-20	400	0.03
		11	0	2	7	-20	400	0.03
990407	I	22	175	9	1	-25	400	0.03
		22	175	1	9	-24	400	0.04
991110	I	33	175	5	4	-22	400	0.04
		30	175	2	8	-23.1	400	0.04
991111	I	22	175	5	4	-22	400	0.04
		23	175	2	8	-23.1	400	0.04
011114	D	21	125	5	4	-22	400	0.04
		12	125	2	8	-23.1	400	0.04
011126	D	18	125	5	4	-22	400	0.04
		18	125	2	8	-23.1	400	0.04
011205	D	17	125	5	4	-22	400	0.04
		17	125	2	8	-23.1	400	0.04

[a] I: states are drawn independently across periods; D: states are drawn without replacement, starting from a population of 18 balls, six of each type (state).

[b] As discussed in the text, endowment of Notes includes loans to be repaid at the end of the period.

across experiments. Indeed, because preferences may be affected by earnings during the experiment, the possibility of bankruptcy, and the time to the end of the experiment, preferences may even be different across periods in the same experiment. Because equilibrium prices and choices depend on all of these, and because of the inevitable noise present in every experiment, there is every reason to expect equilibrium prices and choices to be different across experiments or even across different periods in a given experiment.

Most of the subjects in these experiments had some knowledge of economics in general and of financial economics in particular: Caltech undergraduates had taken a course in introductory finance, Claremont and Occidental undergraduates were taking economics and/or econometrics classes, and MBA students are exposed to various courses in finance. In one experiment (011126), subjects were undergraduates at the University of Sofia (Bulgaria), and were perhaps less knowledgeable about economics and finance.

3 Findings

Because all trading is done through a computerized continuous double auction, we can observe and record every transaction — indeed, every offer — but we focus on end-of-period prices: that is, the prices of the last transaction in each period.[3] Because no uncertainty is resolved while the market is open, it is natural to organize the data using a static model of asset trading: investors trade assets before the state of nature is known, assets yield dividends and consumption takes place after the state of nature is revealed (see Arrow & Hahn (1971) or Radner (1972)).[4]

Because Notes and Cash are both riskless, we simplify slightly and treat them as redundant assets.[5] We therefore model our environment as involving trade in risky assets A, B and a one riskless asset N (notes). Assets are claims to consumption in each of the three possible states of nature X, Y, Z. Write $\operatorname{div} A$ for the state-dependent dividends of asset A, $\operatorname{div} A(s)$ for dividends in state s, and so forth. If $\theta = (\theta_A, \theta_B, \theta_N) \in \mathbb{R}^3$ is a *portfolio* of assets, we write

$$\operatorname{div} \theta = \theta_A(\operatorname{div} A) + \theta_B(\operatorname{div} B) + \theta_N(\operatorname{div} N)$$

for the state-dependent dividends on the portfolio θ.

There are I investors, each characterized by an endowment portfolio $\omega^i = (\omega^i_A, \omega^i_B, \omega^i_N) \in \mathbb{R}^2_+ \times \mathbb{R}$ of risky and riskless assets, and a strictly concave, strictly monotone utility function $U^i : \mathbb{R}^3 \to \mathbb{R}$ defined over state-dependent terminal consumptions. (To be consistent with our experimental design, we allow consumption to be negative but we require holdings of A, B to be non-negative.) Investors care only about consumption, so given asset prices q, investor i chooses a portfolio θ^i to maximize $\operatorname{div} \theta^i$ subject to the budget

[3]See Asparouhova, Bossaerts & Plott (2003) and Bossaerts & Plott (2004) for discussion of the evolution of prices during the experiment.

[4]Because there is only one good, there is no trade in commodities, hence no trade after the state of nature is revealed.

[5]In fact, Cash and Notes are not quite perfect substitutes because all transactions must take place through Cash, so that there is a transaction value to Cash. As Table 3 shows, however, Cash and Notes are nearly perfect substitutes at the ends of most periods in most experiments.

constraint $q \cdot \theta^i \leq q \cdot \omega^i$.

An *equilibrium* consists of asset prices $q \in \mathbb{R}^3_{++}$ and portfolio choices $\theta^i \in \mathbb{R}^2_+ \times \mathbb{R}$ for each investor such that

- choices are budget feasible: for each i

$$q \cdot \theta^i \leq q \cdot \omega^i$$

- choices are budget optimal: for each i

$$\varphi \in \mathbb{R}^2_+ \times \mathbb{R}, U^i(\operatorname{div} \varphi) > U^i(\operatorname{div} \theta^i) \Rightarrow q \cdot \varphi > q \cdot \omega^i$$

- asset markets clear:

$$\sum_{i=1}^{I} \theta^i = \sum_{i=1}^{I} \omega^i$$

In the following subsections, we show first, that observed prices are generally below risk neutral prices, which implies risk aversion; second, that risk aversion is systematic; third that the effects of risk aversion can be quantified; and fourth, that risk aversion can be estimated.

3.1 Risk Neutral Pricing and Observed Pricing

Risk neutrality for investor i means that $U^i(x) = E(x)$ (where the expectation is taken with respect to the true probabilities. If all investors are risk neutral then (normalizing so that the price of Cash is 1 and the price of Notes is 100), the unique equilibrium price is the risk-neutral price $q = (E(A), E(B), E(N)) = (E(A), E(B), 100)$.

Table 3 displays end-of-period prices in 72 periods across 9 experiments: the end-of-period price of asset A is below its expectation in 64 periods, equal to its expectation in 5 periods, above its expectation in 3 periods; the end-of-period price of asset B is below its expectation in 64 periods, equal to its expectation in 3 periods, above its expectation in 5 periods.

Table 3: End-Of-Period Transaction Prices

Date	Sec[a]	Period								
		1	2	3	4	5	6	7	8	9
981007	A	220/230[b]	216/230	215/230	218/230	208/230	205/230			
	B	194/200	197/200	192/200	192/200	193/200	195/200			
	N[c]	95[d]	98	99	97	99	99			
981116	A	215[e]	203	210	211	185	201			
	B	187	194	195	193	190	185			
	N	99	100	98	100	100	99			
990211	A	219	230	220	201	219	230	240		
	B	190	183	187	175	190	180	200		
	N	96	95	95	98	96	99	97		
990407	A	224	210	205	200	201	213	201	208	
	B	195	198	203	209	215	200	204	220	
	N	99	99	100	99	99	99	99	99	
991110	A	203	212	214	214	210	204			
	B	166	172	180	190	192	189			
	N	96	97	97	99	98	101			
991111	A	225	217	225	224	230	233	215	209	
	B	196	200	181	184	187	188	188	190	
	N	99	99	99	99	99	99	99	99	
011114	A	230/230	207/225	200/215	210/219	223/223	226/228	233/234	246/242	209/228
	B	189/200	197/203	197/204	200/207	189/204	203/208	211/212	198/208	203/210
	N	99	99	99	99	99	99	99	98	99
011126	A	180/230	175/222	195/226	183/217	200/220	189/225	177/213	190/219	
	B	144/200	190/201	178/198	178/198	190/201	184/197	188/198	175/193	
	N	93	110	99	100	98	99	102	99	
011205	A	213/230	212/235	228/240	205/231	207/237	232/242	242/248	255/257	229/246
	B	195/200	180/197	177/194	180/194	172/190	180/192	190/195	185/190	185/190
	N	99	100	99	99	99	99	99	99	100

[a]Security.

[b]End-of-period transaction price/expected payoff.

[c]Notes.

[d]For Notes, end-of-period transaction prices only are displayed. Payoff equals 100.

[e]End-of-period transaction prices only are displayed. Expected payoffs are as in 981007. Same for 990211, 990407, 991110 and 991111.

Indeed, in many experiments, all or nearly all transactions take place at a price below the asset expectation. For example, Figure 1 records *all* the purchases/sales of assets throughout the 8 periods of an experiment conducted on November 26, 2001: *all* of the more than 500 trades of the risky assets take place at a price *below* the assets' expected payoffs.

3.2 Prices and Betas

Subsection 3.1 shows that asset prices are below risk neutral prices, which implies risk aversion on the part of subjects. To see that the effect of risk aversion is systematic, we examine expected returns and asset betas.

Recall that the *market portfolio* is the social endowment of all assets

$$M = \sum_{i=1}^{\infty} \omega^i$$

The *beta* of a portfolio θ is the ratio of the covariance of θ with the market portfolio to the variance of the market portfolio

$$\beta(\theta) = \frac{\mathrm{cov}\,(\mathrm{div}\,\theta, \mathrm{div}\,M)}{\mathrm{var}\,(\mathrm{div}\,M)}$$

Given prices q, the *expected rate of return* of a portfolio θ is $E(\mathrm{div}\,\theta/q \cdot \theta)$.

Most asset pricing theories predict that assets with higher betas should have higher expected rates of return. (For example, the Capital Asset Pricing Model predicts $E(\mathrm{div}\,\theta/q \cdot \theta) - 1 = \beta(\theta)\,[E(\mathrm{div}\,M/q \cdot M) - 1]$.) In our laboratory markets, asset A always has higher beta than asset B so should have higher expected rated of return. Figure 2 plots the difference in expected rates of return (expected rate of return of A minus expected rate of return of B) against the difference in betas (beta of A minus beta of B) for all 67 observations (all periods of all experiments). As the reader can see, the difference in expected rate of return is positive roughly 75% of the time. Applying a binomial test to the data yields a z-score of 8, so the correlation is very unlikely to be accidental.

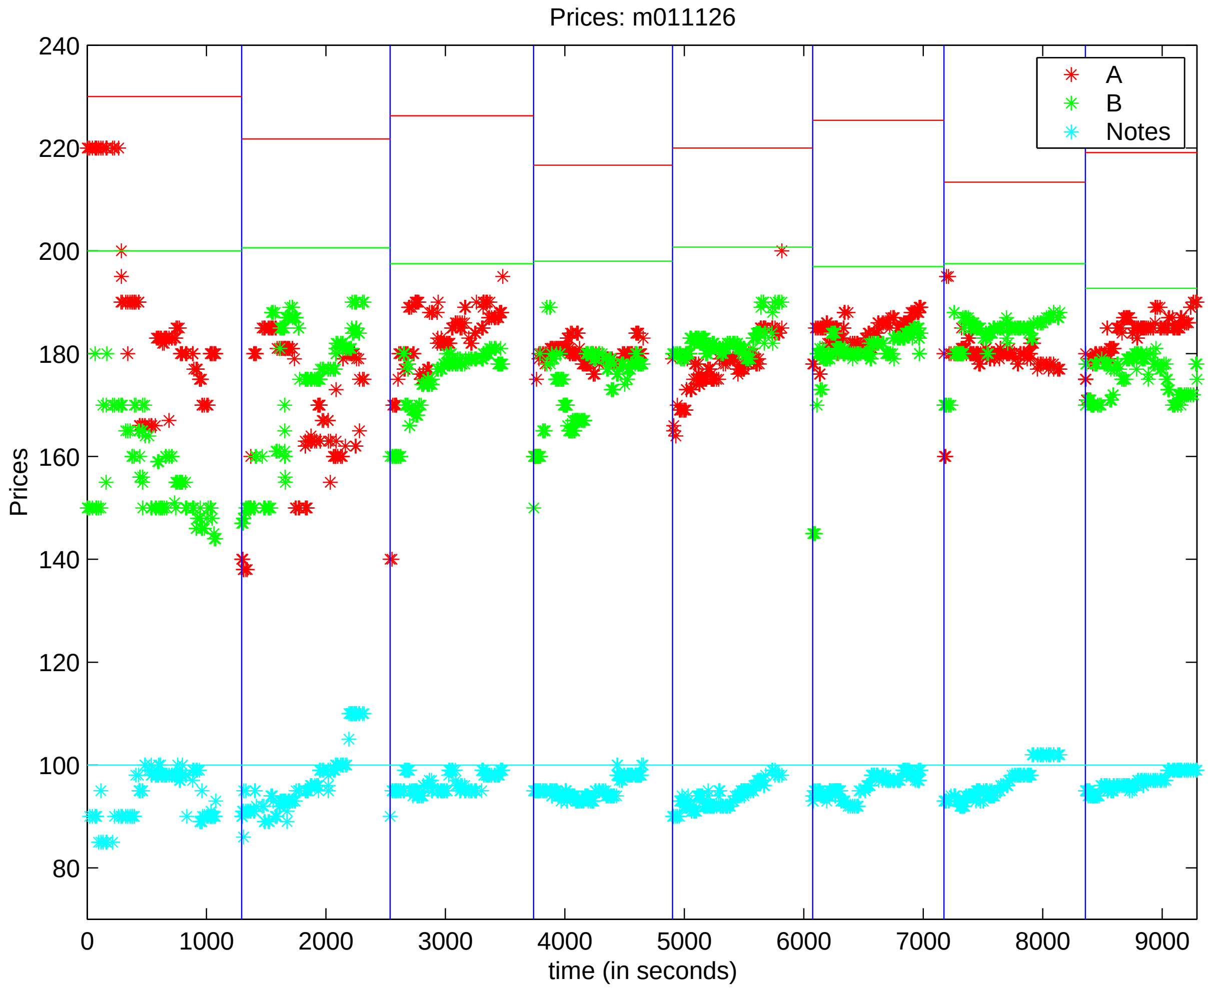

Figure 1: Transaction prices in experiment 011126

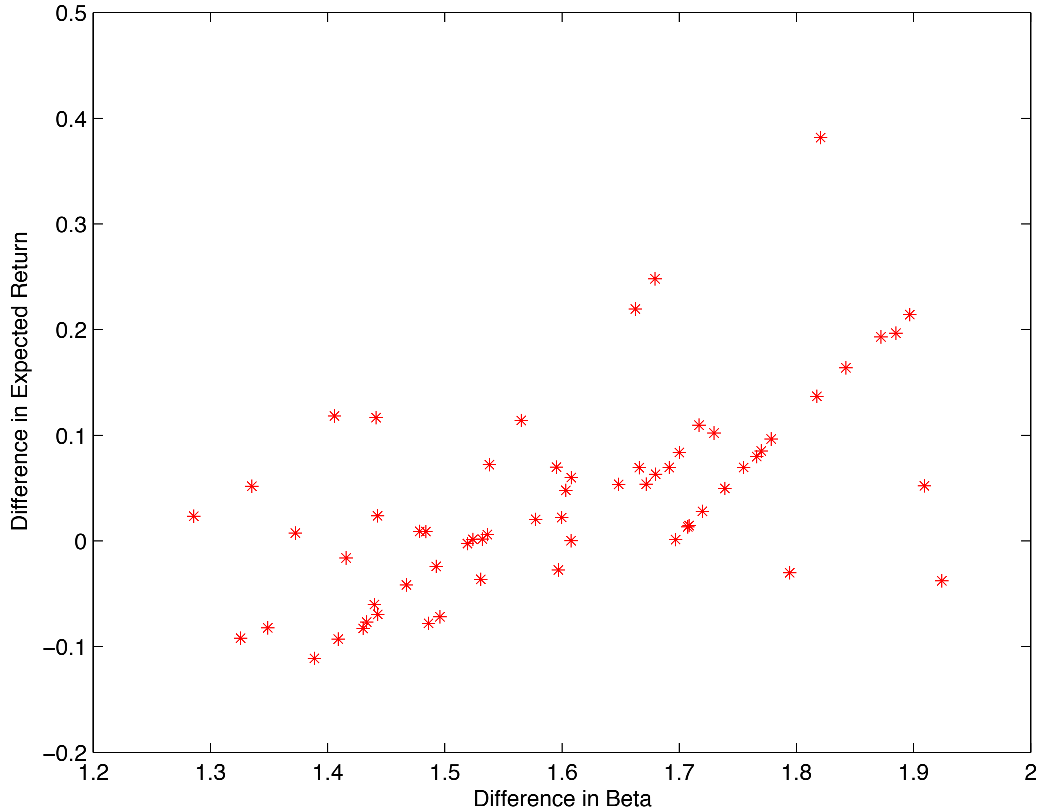

Figure 2: Differences of Betas vs Differences of Expected Returns

3.3 Sharpe Ratios

The data discussed above show that asset prices in our laboratory asset markets reflect significant risk aversion; Sharpe ratios provide a useful way to quantify the effect of this risk aversion. Given asset prices q, the *excess rate of return* is the difference between the rate of return on θ and the rate of return on the riskless asset. In our context, the rate of return on the riskless asset is 1, so the excess rate of return on the portfolio θ is $E[\text{div}\ \theta/q \cdot \theta] - 1$. By definition, the *Sharpe ratio* of θ is the ratio of its excess return to its volatility:

$$\mathrm{Sh}\,(\theta) = \frac{\mathrm{E}[\mathrm{div}\,\theta/\mathrm{q}\cdot\theta] - 1}{\sqrt{\mathrm{var}(\mathrm{div}\,\theta/\mathrm{q}\cdot\theta)}}$$

In particular, the Sharpe ratio of the market portfolio M is

$$\mathrm{Sh}\,(\mathrm{M}) = \frac{\mathrm{E}[\mathrm{div}\,\mathrm{M}/\mathrm{q}\cdot\mathrm{M}] - 1}{\sqrt{\mathrm{var}(\mathrm{div}\,\mathrm{M}/\mathrm{q}\cdot\mathrm{M})}}$$

If investors were risk neutral, asset prices would equal expected dividends, so the numerator would be 0, and the Sharpe ratio of the market portfolio (indeed of every portfolio) would be 0. Roughly speaking, increasing risk aversion leads to lower equilibrium prices and hence to a higher Sharpe ratio (as we see below, CAPM leads to a precise statement), so the Sharpe ratio is a quantitative — although indirect — measure of market risk aversion.

As Figure 3 shows, except for one outlier, Sharpe ratios in our laboratory markets are in the range $0.2 - 1.7$, clustering in the range $0.4 - 0.6$. For comparison, recall that the Sharpe ratio of the market portfolio of stocks traded on the New York Stock Exchange (computed on yearly data) is about .43. (Keep in mind that risks and rewards on the NYSE are enormously greater than in our experiments, so similar Sharpe ratios do not translate precisely into similar risk attitudes.)

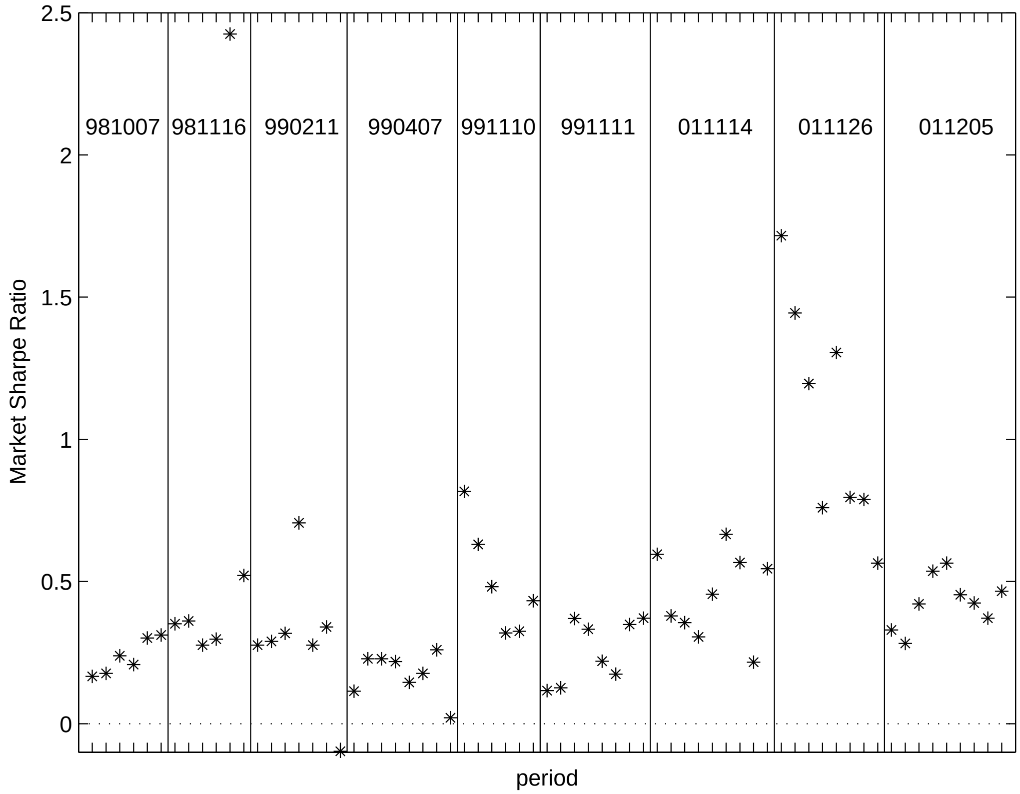

Figure 3: Sharpe Ratios: All Periods, All Experiments

3.4 CAPM

An alternative approach to quantifying the risk aversion in our laboratory markets is to use a particular asset pricing model to impute the market risk aversion. The Capital Asset Pricing Model (CAPM) of Sharpe (1964) is particularly well-suited to this exercise.

CAPM can be derived from various sets of assumptions on primitives. For our purposes, assume that each investor's utility for risky consumption depends only on the mean and variance; specifically, investor i's utility function for state-dependent wealth x is

$$U^i(x) = E(x) - \frac{b^i}{2}\mathrm{var}\,(x)$$

where expectations and variances are computed with respect to the true probabilities, and b^i is absolute risk aversion. We assume throughout that risk aversion is sufficiently small that the utility functions U^i are strictly monotone in the range of feasible consumptions, or at least observed consumptions. Because we allow consumption to be negative, and individual endowments are portfolios of assets, this is enough to imply that CAPM holds.[6]

To formulate the pricing conclusion of CAPM, write $m = \sum(\omega^i_A, \omega^i_B)$ for the *market portfolio of risky assets*, and $\overline{m} = m/I$ for the *per capital portfolio of risky assets*. Write $\mu = (E(A), E(B))$ for the vector of expected dividends of risky assets,

$$\Delta = \begin{pmatrix} \mathrm{cov}\,[A, A]\} & \mathrm{cov}\,[A, B]\} \\ \mathrm{cov}\,[B, A]\} & \mathrm{cov}\,[B, B]\} \end{pmatrix}$$

for the covariance matrix of risky assets, and

$$\Gamma = \left(\frac{1}{I}\sum_{i=1}^{I}\frac{1}{b^i}\right)^{-1}$$

[6] In the usual CAPM, all assets can be sold short, while in our framework the risky assets A, B cannot be sold short. However, in Appendix A of **?** we show that, given the particular asset structure here, the restriction on short sales does not change the conclusions.

for the *market risk aversion*. Write $p = (p_A, p_B)$ for the vector of prices of risky assets. The pricing conclusion of CAPM is that the equilibrium price of risky assets is given by the formula

$$\tilde{p} = \mu - \Gamma\,\Delta\,\overline{m}$$

In our setting, we know equilibrium prices, expected dividends, asset dividends and true probabilities, hence the covariance matrix, and the per capita market portfolio but not individual risk aversiions. If CAPM pricing held exactly, we could impute the market risk aversion by solving the pricing formula for Γ. In our experiments, CAPM pricing does not hold exactly (see Bossaerts, Plott & Zame (2005) for discussion of the distance of actual pricing to CAPM pricing), but we can impute market risk aversion as the best-fitting Γ. Several possible notions of "best-fitting" might be natural; we use Generalized Least Squares, where weights are based on the dispersion of individual holdings from the market portfolio; this is an economic measure of distance used and discussed in more detail in Bossaerts, Plott & Zame (2005). Figure 4 shows the imputed market risk aversion for all periods in all experiments. Note that there is considerable variation across experiments, and even within a given experiment; as we have noted earlier, subject preferences certainly vary across experiments and may even vary within a given experiment.

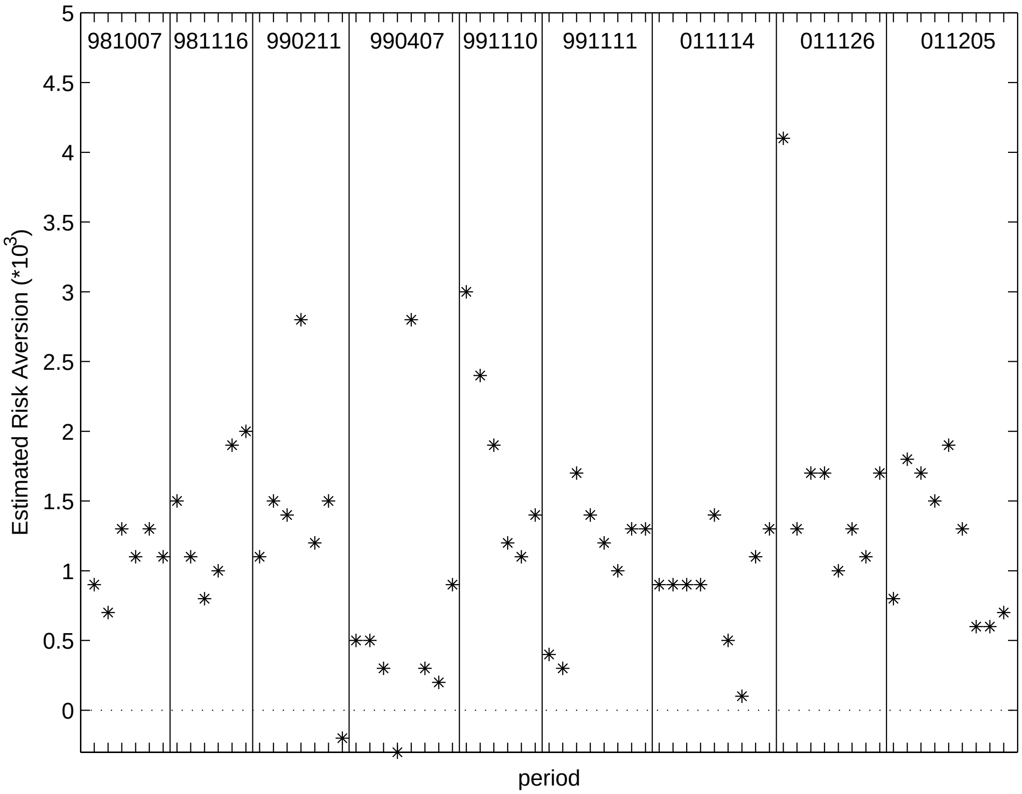

Figure 4: Imputed Market Risk Aversion: All Periods, All Experiments

4 Conclusion

We have argued here that the effects of risk aversion in laboratory asset markets are observable and significant, that the observed effects are in the direction predicted by theory, and that these effects are quantifiable.

A crucial feature of our experimental design is that there are *two* risky assets, so that the realization of uncertainty has two separate, but correlated, effects, and it is this correlation that makes it possible to make quantitative inferences about the effects of risk aversion. This feature suggests an approach to understanding the findings of other laboratory environments in which risk aversion may play a role. For example, in laboratory tests of auction theory, some deviations of observed behavior from theoretical predictions may be interpreted failures of the theory — and hence may point to other theories — or as effects of risk aversion. Our work suggests that these competing explanations might be disentangled by auctioning *two objects* whose values are risky but correlated.

References

KENNETH ARROW AND FRANK HAHN, *General Competitive Analysis*, San Francisco: Holden-Day (1971).

E. ASPAROUHOVA, P. BOSSAERTS AND C. PLOTT, "Excess Demand and Equilibration In Multi-Security Financial Markets: The Empirical Evidence," *Journal of Financial Markets* 6 (2003), 1-22.

FISCHER BLACK AND MYRON SCHOLES, "The Pricing of Options and Corporate Liabilities,' *Journal of Political Economy* 81 (1973), 637-654.

PETER BOSSAERTS AND CHARLES PLOTT, "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," *Review of Finance* 8 (2004), 135-169.

PETER BOSSAERTS, CHARLES PLOTT AND WILLIAM ZAME, "Prices and Portfolio Choices in Financial Markets: Theory and Experiment," Caltech Working Paper (2005).

J. DAVIS, E. FAMA AND K. FRENCH, "Characteristics, Covariances, and Average Returns: 1929 to 1997," *Journal of Finance* 55 (2002), 389-406.

C. HOLT AND S. LAURY, "Risk Aversion and Incentive Effects," *American Economic Review* 92 (2002), 1644-1655.

ROY RADNER, "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica* 40 (1972), 289-303.

WILLIAM SHARPE, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance* 19 (1964), 425-442.