Prices And Portfolio Choices In Financial Markets: Econometric Evidence.

Peter Bossaerts[‡], Charles Plott[§], William Zame[¶]

This version: March 2002 (First version: September 2001)

[‡]California Institute of Technology and CEPR
[§]California Institute of Technology
[¶]UCLA and California Institute of Technology

Prices And Portfolio Choices In Financial Markets:
Econometric Evidence.

Abstract

Structural econometrics is developed to test an extension of standard general equilibrium theory that was suggested in [2] as explanation for why in experiments securities prices tend to the predictions of general equilibrium theory but allocations do not. The test's asymptotic properties are derived as one increases the number of subjects (relatively large in the experiments at hand), keeping fixed the time dimension (small in the experimental setting). The test, based on GMM, is arranged such that errors in the estimation of individual subjects' risk tolerance coefficients average out as the number of subjects increases. Prices are linked to individual portfolio choices, thereby providing a comprehensive test of equilibrium conditions, unlike in studies of field data. The empirical distribution of the test's central statistic is a horizontal translation of its distribution under the null (a χ^2 distribution with one degree of freedom), implying that there are more rejections than expected by chance. Learning (markets move closer to equilibrium in later periods of an experiment) does not explain the finding. Small-sample bias contributes to the rejections, as evidenced by correlation between the test's statistic and a standard measure of the distance from equilibrium pricing. A closer examination of the estimated average allocation prediction errors reveals model mis-specification as an additional source of the rejections: violations of short-sale constraints on the riskfree asset.

Prices And Portfolio Choices In Financial Markets:
Econometric Evidence.

1 Introduction

Consider a set of markets where a large number of agents trade securities which will pay a single, risky liquidating dividend in the near future. There is a riskfree security. The markets are isolated in the sense that agents cannot simultaneously trade in other markets until the dividends are realized, and risk is not correlated with outside risk. In addition, risk is small. It is sufficiently simple, so that the number outcomes ("states") equals the number of nonredundant securities.

Abstracting from the complexities of the actual trading, this case can reasonably be modeled as a one-period set of complete, competitive securities markets. In equilibrium, a representative agent exists. If her preferences are of the expected-utility type and smooth, they could be approximated with a quadratic function, because risk is small. Markets are isolated, so they have their own "market portfolio" (per-capita supply of risky securities). Consequently, one could expect them to generate the CAPM equilibrium.

The CAPM makes the following predictions about prices and allocations. Prices are such that the *market portfolio is mean-variance optimal*, i.e., the difference between its Sharpe ratio (expected excess return divided by volatility) and the maximal Sharpe ratio is zero. Agents will all be holding the same portfolio of risky assets, namely, the market portfolio itself. The latter property is usually referred to as *portfolio separation*.

[1, 2] report the outcomes of a number of financial markets experiments set up as described above, with three securities (denoted A, B, and Notes). Details about the experimental setup can be found in the Appendix. While prices clearly moved in the direction predicted by the CAPM, allocations did not. Table 1 and Figure 1 roughly summarize the results.

Table 1 reports Sharpe ratio differences at the end of each period[1] as well as the maximum difference during the period (in parentheses). The Sharpe ratio differences at the end of the period are often low, and predominantly in the lower half of the interval between zero (the theoretical minimum) and the maximum value for the period (62 out of 83, which corresponds to a p level of 0.01 under the null that the end-of-period Sharpe ratios are equally likely to be in the lower and upper half).

Figure 1, however, demonstrates that there is no relationship between proximity to CAPM pricing (as measured by end-of-period Sharpe ratio differences) and proximity to CAPM allocations. The latter are measured as mean absolute deviations across subjects between the proportion of their risk-exposed wealth allocated to the first of the two risky

[1] Reported are averages based on prices for the last 10 transactions. Each "period" is a replication of the same situation described above. See Appendix.

securities (A) and the corresponding weight in the market portfolio. Under the CAPM, the mean absolute deviations should be zero.

The fact that these mean absolute deviations between actual holdings and CAPM predictions are nonzero may not be surprising, as they may not necessarily entail large utility losses. It is puzzling, however, to observe that allocations are no closer to CAPM when pricing is, for the allocational prediction (portfolio separation) is generally understood to be crucial to obtain CAPM pricing.

Nevertheless, [2] point out that CAPM pricing may still obtain even if the model makes large allocation prediction errors, as long as these errors average out. [2] went on to explore the nature of the allocation prediction errors, attributing the bulk of them to preference heterogeneity (including loss aversion, biased beliefs, non-state-separable preferences, etc.), instead of to inattention, optimization mistakes and trading mistakes. They also estimate the size of the allocation prediction errors and report that they are large, in the sense that they would entail a significant loss in utility if subjects really had mean-variance preferences.

The model in [2], to be referred to henceforth as BPZ model, provides a potential explanation for the absence of correlation depicted in Figure 1, but the link is not tight. The BPZ model makes predictions about allocation prediction errors, defined as deviations between actual holdings of risky securities, on the one hand, and mean-variance optimal holdings given observed prices, on the other hand, and not about deviations in between actual portfolio composition and the composition of the market (plotted in Figure 1). The purpose of this paper is to bring the BPZ model to the data in a structural test built around its core prediction, namely that allocation prediction errors are mean zero.

There is crucial difference between our structural test of (the BPZ version of) CAPM on experimental data and tests used on field data. Ours ties links allocations and prices, unlike in field tests, where generally only pricing is investigated. Note that pricing restrictions are only one aspect of equilibrium. Equilibrium cannot genuinely be tested for without considering allocations, because it could be that CAPM pricing obtains even when demand does not equal supply. Because our tests links prices and allocations, it is a genuine test of equilibration.

Closer inspection of our test statistic will reveal that it effectively measures how far prices are from making the market portfolio mean-variance efficient. That is, it measures the distance of the market portfolio from mean-variance optimality, as in a field test of the CAPM. The standard errors that scale this distance, however, are based on estimates of the cross-sectional variance of allocation prediction errors, unlike in field tests, where standard errors are based on sampling variation in the empirical distribution of future returns. This also means that our test effectively explains deviations from CAPM pricing as sampling variation in allocation prediction errors.

In essence, allocation prediction errors provide the stochastic structure for our econometric test. In field tests, sampling error in the estimated distribution of future returns generates variation in the test statistics. It would be inappropriate to build a structural test for our (experimental) data on the basis of sampling error in the estimated distribution of future payoffs, because the payoff distribution is a control, and hence, known without error.

Our test of the BPZ model is a variation of generalized method of moments (GMM). There is a complication, though: estimates are needed of the risk tolerance (inverse risk aversion parameter) of each subject. Risk tolerances can only be estimated from the time series of end-of-period choices. Because experiments necessarily contain a limited number of periods, the estimation errors of the individual risk tolerances will never disappear, even if we take the cross-section (number of subjects) to be large. By setting up our estimators in a particular way, however, errors in estimating risk tolerances have no asymptotic effect on the GMM χ^2 statistic. Our appealing to large cross-sections to derive the statistical properties of the GMM estimator and test is in line with the nature of the experiments. They involve an unusally large number of subjects (unusual for experimental economics, that is).

The remainder of this paper is organized as follows. The next section briefly describes the nature of the data that the experiments in [2] generate. Section 3 presents the BPZ model. Section 4 derives the corresponding GMM structural test. The empirical results are reported in Section 5. Section 6 concludes. Description of the experiments and technical details of the proofs are collected the Appendix.

2 Nature Of The Data

Our experiments involve multiple markets in simple assets. Without going into detail (the Appendix provides a more thorough description), the experiments are set up as replications, each replication being referred to as a period, of markets in a number of short-lived securities. With the exception of one security, their end-of-period liquidating dividend is stochastic. The actual dividend is determined by a random drawing of a state. Subjects know the nature of the randomness as well as the payoff matrix (mapping from states to dividends). The number of assets equals at least the number of states, and the payoff matrix is full rank, implying that markets are complete. No subject is given superior information.

Subjects are allocated a specific number of each of the assets at the beginning of a period. For up to 25 minutes, subjects can buy and/or sell assets in a set of parallel, continuous open book markets, organized over the web (the mechanism is referred to as Marketscape). All payments are effected using an artificial currency referred to as francs (to avoid having to deal with noninteger prices), converted to U.S. dollars at a pre-announced exchange rate at the end of the experiment. Subjects are given a certain amount of cash at the beginning of each period. They can generate additional cash by selling or (to a certain extent) borrowing. Borrowing takes place through shortsales of the riskfree securities referred to as Notes. At the end of each period, a fixed, pre-announced amount is subtracted from subjects' cash holdings plus the dividend payments on the assets they still hold. This way, subjects can actually have negative earnings in a period. If they generate negative *cumulative* earnings more than two periods in a row, subjects are excluded from further participation (in which case they forego the opportunity to make more money).

Subjects are not allowed to shortsell risky assets, a constraint that is not binding in CAPM and even more general versions of competitive equilibrium, where everybody will end up with the same positively weighted portfolio of risky assets, as we will explain shortly. Shortsales of riskfree securities (Notes) are allowed. Because Notes were always in zero net supply, such shortsales are necessary for efficient risk sharing between subjects with different levels of risk aversion. Nevertheless, subjects could sell short riskfree securities only up to about one-half the value of their initial holdings of risky securities.

The experiments generate a rich dataset that includes bids and bid quantities, asks and ask quantities (at various levels in the book), transaction prices and quantities, and changes in portfolio holdings, all time stamped to the second. Because the scale of the experiments is rather large (relative to standard economics experiments – up to 63 subjects), the data are high frequency and the one-second time stamp is not a luxury. A typical experiment easily generates over 1,000 transactions, and several times that number in terms of bids and asks. Changes in portfolio composition through trading can be accurately followed. For the purpose of this paper, we are going to retain only a subset of the data, namely, end-of-period portfolio holdings and transaction prices. We then ask to what extent the holdings and prices in this subset are arranged in the way predicted by general equilibrium theory.

3 The BPZ Model

Abstracting from the complex, continuous trading in the experiments, we choose to bring general equilibrium theory to bear on the data. As argued in the Introduction, features of the experiments allow us to narrow the analysis down to the CAPM. This model makes a sharp pricing prediction: securities should become priced such that the market portfolio is mean-variance optimal. This comes about because of an extreme feature of demands in the context of quadratic utility (the type of preferences underlying CAPM), namely, portfolio separation: no matter how risk averse a subject is, the optimal demand can always be decomposed into two parts: the riskfree asset and a single benchmark portfolio of risky securities. In equilibrium, this single benchmark portfolio must be the market portfolio.

Inspired by standard practice in applied economics, we now enrich the basic CAPM and add an error term to individual demands. This error term will be referred to as *allocation prediction error*. It is meant to capture preference heterogeneity that is not modeled in the traditional CAPM (which assumes quadratic utility). The allocation prediction error is assumed to be mean zero and independent across subjects.

More specifically, let $h_n^0 \in R_+$ denote subject n's endowment of the riskless asset and $z_n^0 \in \mathbf{R}_+^J$ subject n's endowment of the risky assets. We assum that subject n's demand can be written as the sum of the optimal demand for some quadratic utility function (with parameters that may differ across n) plus a mean-zero, independent error term.

It should be emphasized that our introducing additive error terms at the demand level is a deliberate modeling choice (we could have introduced them at the utility level, for instance). It ensures that standard pricing results (CAPM pricing) are not invalidated, at least as long as there is a large number of subjects, even if portfolio choices do not exhibit standard properties (portfolio separation). We demonstrate this later on in this section.

To fix ideas, we first derive the optimal demands and corresponding equilibrium as if all subjects had quadratic utility. We then add the error terms, to investigate the impact on equilibrium prices and allocations.

Instead of using quadratic utility, we will work with the more popular and equivalent representation in terms of a trade-off between mean return against variance. Write $D_j(s)$ for the return of the j-th risky asset in state s. Let μ be the vector of expected payoffs of risky assets and $\Delta = [\mathrm{cov}\,(D_j, D_k)]$ be the covariance matrix. A subject who holds h units of the riskless asset and the vector z of risky assets will enjoy utility

$$u_n(h, z) = h + [z'\mu] - \frac{b_n}{2}[z'\Delta z] \tag{1}$$

b_n will be referred to as the risk aversion coefficient; its inverse will be called risk tolerance.

Using the price of the riskfree asset as numéraire, let p denote the vector of prices of risky assets. Given prices p, investor n's budget set consists of portfolios (h, z) that yield non-negative consumption in each state and satisfy the budget constraint

$$h + p \cdot z \leq h_n^0 + p \cdot z_n^0 \tag{2}$$

Assuming the portfolio choice yields consumption between 0 and $1/b_n$ in each state of the world (an assumption that held in all our experiments), the first order conditions characterize optima, so investor n's demand for risky assets given prices p is

$$\tilde{z}_n(p) = \frac{1}{b_n}\Delta^{-1}(\mu - p) \tag{3}$$

In particular, all investors choose a linear combination of the riskless asset and the *same* portfolio of risky assets — a conclusion usually referred to as *portfolio separation*.

As usual, an *equilibrium* consists of prices $\tilde{p}$ for assets and portfolio choices $\tilde{h}_n, \tilde{z}_n$ for each investor so that subjects optimize in their budget sets and markets clear. Given the nature of demands (3), market clearing requires that the portfolio $\Delta^{-1}(\mu - \tilde{p})$ be a multiple of the market portfolio $z^0 = \sum z_n^0$ (the supply of risky assets). Solving for equilibrium prices yields

$$\tilde{p} = \mu - \left(\sum_{n=1}^{N} \frac{1}{b_n}\right)^{-1} (\Delta z^0)$$

It is convenient to write $B = (\frac{1}{N}\sum_{n=1}^{N}\frac{1}{b_n})^{-1}$ for the harmonic mean risk aversion and $\bar{z} = \frac{1}{N}\sum_{n=1}^{N} z_n^0$ for the mean endowment (mean market portfolio). With this notation, the pricing formula becomes

$$\tilde{p} = \mu - B\Delta\bar{z} \tag{4}$$

This is the *Capital Asset Pricing Model (CAPM).*

We now extend the CAPM such that actual portfolio choice is represented as a perturbation of the theoretical (CAPM–optimal) portfolio choice. Formally, we suppose that actual demand z_n is the sum of the theoretical demand $\tilde{z}_n$ and an error ϵ_n, which, as mentioned before, we refer to as allocation prediction error:

$$z_n = \tilde{z}_n(p) + \epsilon_n \tag{5}$$

We *do not* assume allocation prediction errors are small. That is, ϵ_n is not to be interpreted as an approximation error. We only impose that allocation prediction errors are independent across n, with mean zero

$$E[\epsilon_n] = 0, \tag{6}$$

and finite variance V

$$E[\epsilon_n^2] = V > 0. \tag{7}$$

Direct computation shows that equilibrium prices will now be:

$$p = \mu - B\Delta z^0 + B\Delta \frac{1}{N}\sum_{n=1}^{N} \epsilon_n \tag{8}$$

Hence the difference between equilibrium prices p and CAPM prices $\tilde{p}$ is

$$p - \tilde{p} = B\Delta \frac{1}{N}\sum_{n=1}^{N} \epsilon_n \tag{9}$$

Note that the difference between equilibrium prices and CAPM pricing is proportional to the average error term:

$$\frac{1}{N}\sum_{n=1}^{N} \epsilon_n.$$

Sampling variation in this average will generate fluctuations of equilibrium prices around CAPM pricing. On average, CAPM pricing obtains. This is the model that we would like to confront with the experimental data.

Before we do so, notice the effect of adding subjects, ensuring that the endowment average z^0 remains finite, risk tolerances $1/b_n$ have a finite non-zero mean, and perturbations ϵ_n continue to be independent with zero mean and finite variance. To emphasize the dependence on the number of subjects, write $\tilde{p}^N, p^N$ for equilibrium prices in the CAPM economy and perturbed economy with N investors. The Law of Large Numbers implies that, as $N \to \infty$

$$|p^N - \tilde{p}^N| \quad \to \quad 0 \;\text{ a.s.}, \tag{10}$$

$$\frac{1}{N}\sum_{n=1}^{N} [z_n - \tilde{z}_n]^2 \quad \to \quad V \;\text{ a.s..} \tag{11}$$

That is, for large N equilibrium prices in the true economy will (with high probability) be close to equilibrium prices in the benchmark CAPM economy, but equilibrium portfolio choices in the true economy will *not* be close to equilibrium portfolio prices in the benchmark CAPM economy.

4 The GMM Structural Test

Our model provides a unified framework to develop a structural test on experimental data. Typically, we have data from several periods, each period being a replication of the same environment. We'd like to test our theory on the prices and allocations at the end of each period. We present an estimation and testing strategy within the framework of generalized method of moments (GMM).

4.1 The GMM Estimator

The idea may seem straightforward, but some of the details are not, mainly because we insist that the estimators exhibit standard properties as *only* the cross-sectional size (number of subjects) increases, keeping the time series length (number of periods) *constant*. Indeed, while our trading mechanism (Marketscape) opens up the possibility of experiments with very large subject numbers, one is still limited in the number of periods that could practically be squeezed into an experimental session.

We begin with introducing a time index, to reflect the nature of the data, which consists of prices and portfolio choices in each of T periods. Use t for this time index: $t = 1, 2, ..., T$. As before, $n = 1, ..., N$ is the index for the subjects. This way, p_t^N denotes the (vector) of prices of risky securities in period t given N subjects. Although we could have chosen differently, we take the end-of-period transaction prices. Likewise, $z_{n,t}$ denotes subject n's holdings of risky securities at the end of period t. In addition, $\tilde{z}_n(p_t^N)$ denotes the optimal choice for someone exhibiting mean-variance risk attitude with coefficient of risk aversion that matches subject n's choices most closely. It will become clear what we mean with closeness when we present our estimator of individual risk aversion coefficients (actually, risk tolerances, which are the inverse of risk aversion coefficients). Choices $\tilde{z}_n(p_t^N)$ depend on prices as follows:

$$\tilde{z}_n(p_t^N) = \frac{1}{b_n}\Delta^{-1}(\mu - p_t^N).$$

The difference between $z_{n,t}$ and $\tilde{z}_n(p_t^N)$ is the allocation prediction error $\epsilon_{n,t}$:

$$\epsilon_{n,t} = z_{n,t} - \tilde{z}_n(p_t^N).$$

We impose the following restriction on the allocation prediction error.

$$E[\epsilon_{n,t}|p_t^N] = 0, \tag{12}$$

with $\epsilon_{n,t}$ mutually independent across n given p_t^N. We also assume that the conditional variances of $\epsilon_{n,t}$ are finite, and impose an orthogonality condition over time:

$$E[\epsilon_{n,t}\epsilon_{n,s}|p_t^N, p_s^N] = 0, \tag{13}$$

for $t, s = 1, 2, ..., T$, $t \neq s$.

Unlike in the equivalent assumptions in the theory of the previous section, we condition explicitly on prices. This means that we treat prices as fixed across replications of our experiments in the derivation of the statistical properties of our test. A substantial simplification ensues.

Such conditioning does imply that one major difference between the asset pricing model and the econometric model. In the former, we impose eqn. (6). Since prices obviously correlate with the allocation prediction errors $\epsilon_{n,t}$, through the equilibrium condition in (8), eqn. (12) does not follow. But this is the key assumption in our econometric model.

Nevertheless, the two models converge asymptotically, at a speed that is sufficient not to affect the asymptotic distributional characteristics of the GMM test. More specifically, our asset pricing model implies

$$\lim_{N\to\infty} E[\epsilon_{n,t}|p_t^N] = 0.$$

That is, the restriction in our econometric model obtains asymptotically. To see this, note that, as N increases, prices converge to the traditional CAPM (see (10)), where they are obviously independent of individual allocation prediction errors $\epsilon_{n,t}$. We also ought to prove that our asset pricing model converges to the econometric model at the right speed. In particular,

$$\lim_{N\to\infty} \sqrt{N} E[\epsilon_{n,t}|p_t^N] = 0. \tag{14}$$

To prove that this obtains, we would have to show that

$$\sqrt{N} E[(\epsilon_{n,t}) h(p_t^N)'] \to 0,$$

for any measurable function h. Let us limit our attention to demonstrating that it is true for linear functions $h(p_t^N)$ $= \Theta + \Gamma p_t^N$. From (8),

$$\begin{aligned} &\sqrt{N} E[\epsilon_{n,t}(\Theta + \Gamma p_t^N)'] \\ &= \sqrt{N} E[\epsilon_{n,t} \left(\Theta + \Gamma \tilde{p}^N + B\Gamma\Delta \frac{1}{N} \sum_{m=1}^{N} \epsilon_{m,t} \right)'] \\ &= \frac{1}{\sqrt{N}} BE[\epsilon_{n,t}(\epsilon_{n,t})' \Delta \Gamma'] \\ &\to 0. \end{aligned}$$

In principle, we could formulate our econometric model to be identical to the asset pricing model, and use (14). But this would complicate the analysis without adding any new econometric insights. Indeed, the arguments would be standard, having been appealed to many times before, among others in the context of Pitman drift.[2]

[2]The interested reader could consult, e.g., [3], p. 154 ff.

We are ready to propose our GMM estimator. The estimator provides a consistent estimate of the harmonic mean risk aversion B^N (the superscript N is added to reflects its dependence on the sample size).

Remember:

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

Now define $g_{n,t}$:

$$g_{n,t} = B^N \epsilon_{n,t}. \tag{15}$$

From (12), it follows that

$$E[g_{n,t}|p_t^N] = b^N E[\epsilon_{n,t}|p_t^N] = 0. \tag{16}$$

This restriction suggests the following GMM estimation strategy. Define:

$$h_{N,t} = \frac{1}{N} \sum_{n=1}^{N} g_{n,t}. \tag{17}$$

Now estimate the parameter B^N by minimizing the following quadratic form:

$$\min_{B^N} [\sqrt{N} h'_{N,t}] A^{-1} [\sqrt{N} h_{N,t}], \tag{18}$$

where the weighting matrix A is a consistent estimate of the asymptotic covariance matrix of $\sqrt{N} h_{N,t}$. (We will discuss shortly how to obtain A.)

Notice that, given A, the criterion function does not depend on the N individual risk aversion parameters b_n. To see this, write $h_{N,t}$ explicitly as a function of prices and choices:

$$\begin{aligned} h_{N,t} &= \frac{1}{N} \sum_{n=1}^{N} g_{n,t} \\ &= \frac{1}{N} \sum_{n=1}^{N} B^N \epsilon_{n,t} \\ &= \frac{1}{N} \sum_{n=1}^{N} \left(B^N z_{n,t} - \frac{B^N}{b_n} \Delta^{-1} (\mu - p_t) \right) \\ &= B^N \frac{1}{N} \sum_{n=1}^{N} z_{n,t} - \Delta^{-1} (\mu - p_t). \end{aligned} \tag{19}$$

The corresponding GMM test verifies whether the criterion function in (18) is significantly different from zero. Because the weighting matrix A must be positive definite, GMM tests whether the vector in (19) is zero, making it effectively a test of mean-variance optimality of the market portfolio. This is because $\frac{1}{N} \sum_{n=1}^{N} z_{n,t}$, the average holdings of risky securities, is the market portfolio, by definition, and if we set (19) equal to zero, we obtain the optimality condition for risk aversion parameter B^N (compare to (3)).

That is, the GMM test effectively verifies mean-variance optimality of the market portfolio, as if it were a test on field data. Unlike in the latter, however, distance is not measured in function of error in the estimation of the distribution of payoffs. Instead, it is measured in terms of the variances (and covariances) of the allocation prediction errors of the traditional CAPM, i.e., the weighting matrix A. Consequently, both price and allocation information is used to test the BPZ extension to the CAPM.

Our GMM estimator will have the usual asymptotic properties. Letting $\hat{B}^N$ denote the GMM estimator of B^N, this means:

$$|\hat{B}^N - B^N| \to 0,$$

almost surely, and, when evaluated at $\hat{B}^N$,

$$[\sqrt{N} h'_{N,t}] A^{-1} [\sqrt{N} h_{N,t}]$$

is asymptotically χ^2 distributed, with degrees of freedom equal to the length of the vectors $h_{N,t}$ (i.e., the number of risky securities) minus 1.

These properties will obtain only if the weighting matrix A is a consistent estimator of the asymptotic covariance matrix of $\sqrt{N} h_{N,t}$. It is not obvious which estimator will work. Let us now discuss one possible procedure.

4.2 An Estimator Of The Weighting Matrix A

At the outset, one should realize that the usual estimator will not work. Because of our assumption that the allocation prediction errors $\epsilon_{n,t}$ and hence, the $g_{n,t}$s are independent, A would traditionally be obtained as the sum of the sample covariance matrices of the $g_{n,t}$s when evaluated at the parameter value that minimizes the GMM criterion function for $A = I$ (the identity matrix). The problem is that the sample covariance of each $g_{n,t}$ depends on its own incidental parameter b_n.

To obtain A, estimates of the individual risk aversion coefficients or their inverse, the risk tolerances, are necessary. We generate estimates of the individual risk tolerances from the time series of individual portfolio choices. Since we wish to keep the time series length fixed and small, estimation errors are inevitable and potentially large. We write A as a function of the estimates of the risk tolerances in a way that affords large estimation errors that need not disappear as we increase the number of subjects. In particular, we let the estimates enter A linearly and ensure that they are multiplied at most with the allocation prediction error $\epsilon_{n,t}$. We choose an unbiased estimator with an estimation error that is uncorrelated with $\epsilon_{n,t}$. Linearity, unbiasedness and lack of correlation with $\epsilon_{n,t}$ altogether imply that the influence on A of the estimation errors in the risk tolerances disappears asymptotically, by the law of large numbers.

Let us first discuss estimation of the risk tolerance parameters. Consider the time series history of subject n's end-of-period holdings:

$$z_{n,t} = \tilde{z}_n(p_t^N) + \epsilon_{n,t} \tag{20}$$
$$= \frac{1}{b_n}\Delta^{-1}(\mu - p_t^N) + \epsilon_{n,t} \tag{21}$$

The assumption in (12) ($E[\epsilon_{n,t}|p_t^N] = 0$) implies that the regression function of $z_{n,t}$ onto $\Delta^{-1}(\mu - p_t^N)$ is linear. The slope of each of the projections of the elements of $z_{n,t}$ onto the corresponding elements of $\Delta^{-1}p_t^N$ equals the risk tolerance. Hence, the ordinary least squares (OLS) estimate of the slope is an unbiased estimate of the risk tolerance. There will be several estimates: one for each element of $z_{n,t}$. We aggregate them simply by taking the average. The resulting estimate of the risk tolerance remains unbiased. We repeat this for all subjects, generating N estimated risk tolerances.

When computing the weighting matrix A for period t, we estimate the risk tolerances on the basis of the history of subjects' holdings at the end of all periods *except* period t. Our deletion of period-t holdings in the estimation of the risk tolerances is deliberate. Let $\beta_{n,t}$ denote our estimate of subject n's risk tolerance. The estimation error, $\beta_{n,t} - \frac{1}{b_n}$, depends linearly on the allocation prediction errors $\epsilon_{n,\tau}$ for all periods τ except $\tau = t$.[3] The assumption in (13) implies that $\beta_{n,t} - \frac{1}{b_n}$ is uncorrelated with $\epsilon_{n,t}$.[4] Later on, it will become clear that this ensures that the effect

[3] Applying OLS analysis, define

$$\beta_{n,t}^j = \frac{\mathrm{cov}([z_{n,\tau}]_j, [\Delta^{-1}(\mu - p_\tau^N)]_j)}{\mathrm{var}([\Delta^{-1}(\mu - p_t^N)]_j)},$$

where cov and var denote the sample covariance and variance, respectively, over τ in $1, ..., T$ with $\tau \neq t$. Also, $j = 1, ..., J$, with J denoting the number of risky securities (length of the vector $z_{n,\tau}$). $[y]_j$ denotes the jth element of the vector y. With this notation, our estimator of the risk tolerance parameter equals

$$\beta_{n,t} = \frac{1}{J}\sum_j \beta_{n,t}^j.$$

Borrowing further results from OLS analysis, we obtain

$$\beta_{n,t}^j - \frac{1}{b_n} = \frac{\mathrm{cov}([\epsilon_{n,\tau}]_j, [\Delta^{-1}(\mu - p_\tau^N)]_j)}{\mathrm{var}([\Delta^{-1}(\mu - p_t^N)]_j)}.$$

Therefore,

$$\beta_{n,t} - \frac{1}{b_n} = \frac{1}{J}\sum_j \frac{\mathrm{cov}([\epsilon_{n,\tau}]_j, [\Delta^{-1}(\mu - p_\tau^N)]_j)}{\mathrm{var}([\Delta^{-1}(\mu - p_t^N)]_j)}.$$

The sample covariances in the last expression are linear in the allocation prediction errors $[\epsilon_{n,\tau}]_j$. It follows that the estimation error $\beta_{n,t} - \frac{1}{b_n}$ is as well.

[4] To see this, notice:

$$E\{[\epsilon_{n,t}]_k(\beta_{n,t} - \frac{1}{b_n})|p_1^N, ..., p_T^N\} = \frac{1}{J}\sum_j \frac{\mathrm{cov}(E\{[\epsilon_{n,t}]_k[\epsilon_{n,\tau}]_j|p_1^N, ..., p_T^N\}, [\Delta^{-1}(\mu - p_\tau^N)]_j)}{\mathrm{var}([\Delta^{-1}(\mu - p_t^N)]_j)} = 0,$$

for all k ($k = 1, ..., J$).

on A of the errors in the estimation of the individual risk tolerances disappear as N increases.[5]

As mentioned before, A should be a consistent estimator of the asymptotic covariance matrix of $\sqrt{N}h_{N,t}$. We don't use the most straightforward estimator. Instead, we use one where the influence of the estimation errors of the individual risk tolerances averages out in large cross-sections (N large), as already mentioned. Because the $\epsilon_{n,t}$s are independent across n, the asymptotic covariance matrix of $\sqrt{N}h_{N,t}$ equals the following.

$$\begin{aligned}
&\lim_{N\to\infty} NE[h_{N,t}\, h_{N,t}{}'\,|p_t^N] \\
&= \lim_{N\to\infty}(B^N)^2\frac{1}{N}\sum_{n=1}^{N}E[\epsilon_{n,t}\,\epsilon_{n,t}{}'\,|p_t^N] \\
&= \lim_{N\to\infty}(B^N)^2\left\{\frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\,\epsilon_{n,t}{}' - \frac{1}{N}\sum_{n=1}^{N}\left(\epsilon_{n,t}\,\epsilon_{n,t}{}' - E[\epsilon_{n,t}\,\epsilon_{n,t}{}'\,|p_t^N]\right)\right\} \\
&= \lim_{N\to\infty}(B^N)^2\frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\,\epsilon_{n,t}{}' \qquad (22)
\end{aligned}$$

Note that this is not the limit of the sample covariance matrix of holdings across individuals, which we will denote $\mathrm{cov}(z_{n,t})$. Let

$$\bar{z}_N = \frac{1}{N}\sum_{n=1}^{N} z_{n,t}. \qquad (23)$$

$\mathrm{cov}(z_{n,t})$ equals:

$$\mathrm{cov}(z_{n,t}) = \frac{1}{N}\sum_{n=1}^{N}(z_{n,t}-\bar{z}_N)\,(z_{n,t}-\bar{z}_N)'.$$

We propose the following formula for A.

$$\begin{aligned}
A \;=\; &(\hat{B}^N)^2\left\{\mathrm{cov}(z_{n,t}) - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\beta_{n,t}-\frac{1}{N}\sum_{\nu=1}^{N}\beta_{\nu,t})(z_{n,t}-\bar{z}_N)(\mu-p_t^N)'\Delta^{-1}\right. \\
&\left. -\frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\beta_{n,t}-\frac{1}{N}\sum_{\nu=1}^{N}\beta_{\nu,t})\Delta^{-1}(\mu-p_t^N)(z_{n,t}-\bar{z}_N)'\right\}. \qquad (24)
\end{aligned}$$

Let us first demonstrate here that estimation errors of the risk tolerances have no effect on A asymptotically. Consider the second term on the right hand side in (24). The arguments for the third term are analogous. Ignoring the leading factor, we can rewrite it in terms of the true risk tolerance parameters plus estimation errors:

$$\frac{1}{N}\sum_{n=1}^{N}(\beta_{n,t}-\frac{1}{N}\sum_{\nu=1}^{N}\beta_{\nu,t})(z_{n,t}-\bar{z}_N)(\mu-p_t^N)'\Delta^{-1}$$

[5] If we included $\tau = t$, then the estimation error $\beta_{n,t} - \frac{1}{b_n}$ and the allocation prediction error ϵ_t is correlated. But the correlation can be expected to be small and not to affect the GMM estimation much. This was borne out in the results. Still, we only report the results with $\tau \neq t$.

$$
\begin{aligned}
= \quad & \frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\
& + \frac{1}{N}\sum_{n=1}^{N}\left((\beta_{n,t} - \frac{1}{b_n}) - \frac{1}{N}\sum_{\nu=1}^{N}(\beta_{\nu,t} - \frac{1}{b_\nu})\right)(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1}
\end{aligned}
$$

Consider the deviations of portfolio choices from the grand mean in the second term of the last expression, $z_{n,t} - \bar{z}_N$, $n = 1, ..., N$. These depend linearly on the allocation prediction errors $\epsilon_{n,t}$, $n = 1, ..., N$. In the same term, the estimation errors, namely, $\beta_{n,t} - \frac{1}{b_n}$ and $\beta_{\nu,t} - \frac{1}{b_\nu}$, are mean zero, by unbiasedness of the OLS estimator. At the same time, they are uncorrelated with the allocation prediction errors $\epsilon_{n,t}$, because they depend linearly on allocation prediction errors $\epsilon_{n,\tau}$ for $\tau \neq t$, as emphasized before. Clearly, the second term in the above expression is simply the sample covariance of, on the one hand, linear transformations of allocation prediction errors $\epsilon_{n,\tau}$ for $\tau \neq t$, and on the other hand, linear transformations of the allocation prediction errors $\epsilon_{n,t}$. In expectation, this sample covariance is zero. Because allocation prediction errors $\epsilon_{n,\tau}$ and $\epsilon_{n,t}$ are assumed independent across n, the law of large numbers implies that the sample covariance will converge to its expectation. Consequently, the second term in the above expression is zero asymptotically, leaving only the first term. The random behavior of the first term does not depend on errors in the estimation of the risk tolerances. We have the desired result: asymptotically, estimation errors have no impact on A.

It is now straightforward to demonstrate that A provides a consistent estimator of the asymptotic covariance matrix of $\sqrt{N}h_{N,t}$. Because the proof is tedious, we delegate it to the Appendix.

Our choice of A has the main advantage that it can accommodate the errors in estimation of the incidental risk tolerance parameters. They will not affect the results asymptotically. Against that, however, one should hold that that A is not necessarily positive definite in small samples (N small). At this point, it is not obvious how to guarantee positive definiteness.

5 Empirical Results

Let us first recapitulate the results in [2] and summarized in Tables 1 and Figure 1. While prices move towards CAPM by the end of the trading period (Table 1), allocations are not as predicted by CAPM and do not even correlate with distance from CAPM pricing (Figure 1). Distance from CAPM pricing is measured as the difference between the Sharpe ratio (ratio of expected return over volatility) of the market portfolio and the maximal Sharpe ratio.[6] The

[6] In computing the maximal Sharpe ratio, shortsale constraints on Securities A and B were accounted for, except in the case of 001113. The latter experiment involved trade in a complete set of state securities as well as Notes, implying that one of the securities was redundant. We therefore expressed the market portfolio and optimal portfolios in terms of investments in two of the three state securities (X and Y). When expressed as supplies of X and Y only, even the market portfolio involves shorting of X. Therefore, we computed the maximal Sharpe

success of CAPM in predicting allocations is measured as the mean absolute deviation between the proportion of their risk-exposed wealth subjects allocate to security A and the market portfolio weight of A.

The BPZ model was developed as a straightforward extension of general equilibrium theory that would potentially explain this price-allocation paradox. Table 2 reports the results from the GMM goodness-of-fit test of the core prediction of the BPZ model, namely, that the expected allocation prediction error of mean-variance utility theory equals zero. The table leads to the following observations.

First, the estimates of the harmonic mean risk aversion B^N are almost uniformly positive and significant, rejecting risk neutrality. The estimates of B^N are of the same order of magnitude across experiments, although they tend to be higher when the GMM test rejects.

Second, while the p levels of the GMM χ^2 statistics are often acceptable, there are far more rejections than expected by chance. There appears to be some value to the BPZ model, but it is rejected overall. Figure 2 illustrates this aptly. It plots the empirical distribution of the (logarithm of the) χ^2 statistic across the 83 periods of our experiments. The empirical distribution (jagged line) is a horizontal translation of the theoretical distribution of the (logarithm of the) GMM statistic (smooth line). It is as if the shape of the distribution is well specified, but not the location.

One possible explanation for the high number of rejections would be that it takes time for markets to equilibrate, even several periods, but that subjects learn. Under this scenario, numerous rejections in early periods of the experiments are to be expected, with the frequency decreasing over time. Contrary to this conjecture, there is little evidence of a gradual decrease in the χ^2 statistic as a function of period number. See the evidence in Figure 3.[7]

Small-sample biases may account for part of the high frequency of rejections. We derived the properties of the GMM criterion function for large samples only, whereas our samples have as few as 15 observations (in experiment 000804). One small-sample feature that should be immediately apparent is correlation between the GMM statistic and the distance from CAPM pricing. In constructing our GMM test, we ruled out such correlation – see Equation (12), which stated that allocation prediction errors are mean-independent of prices, and hence, uncorrelated with functions of prices such as distance from CAPM pricing.[8] Figure 4 plots the relationship between the (logarithm of) the GMM χ^2 statistic and the distance from CAPM pricing. A positive relationship emerges. The correlation coefficient is estimated to be 0.34, which is significant at the 1% level. The solid line depicts a piecewise linear spline fit. The fit remains high and stays above 1 even when prices are close to CAPM (asymptotically, the GMM χ^2 statistic has an expectation equal to 1, which means that its logarithm has an expectation strictly less than 0).

A pronounced pattern in mean allocation prediction errors generated by the GMM estimation procedure suggests

ratio without shortsale constraints on the risky securities.

[7]Table 1 also reveals that, with the exception of 991110, there is no tendency of pricing to become closer to CAPM over time. As a matter of fact, prices are already close to CAPM at the end of the first period in experiments such as 981116, 990407, 991111 and 011205.

[8]We did emphasize that allocation prediction errors become mean-independent of prices in large samples, so that Equation (12) is consistent with the BPZ model at least asymptotically.

another source for the high frequency of rejections. Estimates of the mean allocation prediction errors are reported in Table 3. With the exception of 001113, the mean allocation prediction error for security A is almost invariably positive, while that for security B is almost always negative.

Notice incidentally that the average allocation prediction errors are generally higher than 1. They are measured in units of each security. The fact that they are generally above 1 therefore indicates that discreteness[9] cannot explain the high frequency of rejections either.[10]

The pattern in mean allocation prediction errors can be explained as the effect of binding shortsale constraints on the riskfree security. Sufficiently risk tolerant subjects would like to borrow more than they are allowed to (8 Notes). Because they are shortsale-constrained, they optimally choose not to buy a mean-variance optimal portfolio. If there are enough shortsale-constrained subjects, this will affect pricing.

To see how binding shortsale constraints affect average allocation prediction errors, consider the following. In the analysis, we will need the value of the entries of a particular vector, namely, $\Delta^{-1}\mu$. In all the experiments except 001113 (the latter is exceptional for other reasons as well, as we point out later), the second entry of this vector (to be referred to below as $[\Delta^{-1}\mu]_2$) is higher than the first one (to be referred to below as $[\Delta^{-1}\mu]_1$) by a factor of 4.

Let there be two types of agents, $n = 1, 2$, exhibiting mean-variance preferences with coefficients $b_1 < b_2$, i.e., type-1 agents are less risk averse. Imagine that, at prices p, type-1 agents are shortsale-constrained, whereas type-2 agents are not. The optimal demands $z_n^C(p)$ are $(n = 1, 2)$:

$$z_n^C(p) = \frac{1}{b_n}\Delta^{-1}(\mu - \lambda_n p),$$

where $\lambda_1 > 1$ (reflecting the shortsale constraint) and $\lambda_2 = 1$ ($i = 2$ is not shortsale-constrained). These demands imply the following equilibrium price vector, denoted p^C. As before, let

$$B^N = \left(\frac{1}{2}\sum_{n=1}^{2}\frac{1}{b_n}\right)^{-1},$$

the harmonic mean of the b_ns. Define

$$B^C = \left(\frac{1}{2}\sum_{n=1}^{2}\frac{1}{b_n/\lambda_n}\right)^{-1}.$$

B^C is the harmonic mean of adjusted coefficients b_n/λ_n. Because $\lambda_1 > 1$ and $\lambda_2 = 1$, $B^N > B^C$, a fact we will exploit later on. The equilibrium price vector equals:

$$p^C = \frac{B^C}{B^N}\mu - B^C\Delta\bar{z}. \tag{25}$$

[9] Subjects could only trade integer quantities of the securities.

[10] Likewise, direct analysis of sensitivity of the GMM χ^2 statistics revealed no noticeable effect from discreteness.

Note that, at prices P^C, the market portfolio will not be mean-variance optimal. That is, the CAPM does not obtain.

Let us now determine the average allocation prediction error that our GMM procedure would estimate in this case. The procedure finds the parameter $\hat{B}^N$ that minimizes the norm of the vector of scaled average allocation prediction errors:

$$\frac{1}{2}\sum_{n=1,2} g_n = \frac{1}{2}\sum_{n=1,2} \hat{B}^N \epsilon_n = \hat{B}^N \frac{1}{2}\sum_{n=1,2} z_n^C(p^C) - \Delta^{-1}(\mu - p^C).$$

Effectively, our procedure searches for the mean-variance efficient portfolio closest to the actual average demands of the agents. Because the average demands of the agents necessarily add up to the market portfolio (i.e., the average supply $\bar{z}$), we can simplify the vector of scaled average allocation prediction errors:

$$\frac{1}{2}\sum_{n=1,2} \hat{B}^N \epsilon_n = \hat{B}^N \bar{z} - \Delta^{-1}(\mu - p^C).$$

Replacing p^C with its equilibrium expression in terms of primitives of the economy, equation (25), and rearranging, we obtain:

$$\frac{1}{2}\sum_{n=1,2} \hat{B}^N \epsilon_n = (\hat{B}^N - B^C)\bar{z} - \frac{B^N - B^C}{B^N}\Delta^{-1}\mu.$$

As norm, GMM uses the quadratic form defined in terms of the symmetric, positive-definite weighting matrix A^{-1}; see (18). The first-order conditions of the GMM minimization problem are:

$$\left((\hat{B}^N - B^C)\bar{z}' - \frac{B^N - B^C}{B^N}\mu'\Delta^{-1} \right) A^{-1}\bar{z} = 0.$$

Straightforward algebra generates an expression for the GMM estimator $\hat{B}^N$:

$$\hat{B}^N = B^C + \frac{1}{\bar{z}'A^{-1}\bar{z}} \frac{B^N - B^C}{B^N}\mu'\Delta^{-1}A^{-1}\bar{z}.$$

Note that $\bar{z}'A^{-1}\bar{z} > 0$ and remember that $B^N > B^C$. Since $\Delta^{-1}\mu > 0$ in all the experiments (except 001113), we conclude that

$$\hat{B}^N > B^C (> 0).$$

All this allows us to obtain conditions on the harmonic mean risk aversion B^N under which the average allocation prediction error of the first security is higher than that of the second security, as is the case in most periods of our experiments (except 001113).

$$\frac{1}{2}\sum_{n=1,2} \epsilon_n = \frac{\hat{B}^N - B^C}{\hat{B}^N} \left(\bar{z} - \frac{B^N - B^C}{\hat{B}^N - B^C} \frac{1}{B^N} \Delta^{-1}\mu \right).$$

Straightforward algebra then shows that the average allocation prediction error of the first security is higher than that of the second security whenever

$$[\bar{z}]_2 - [\bar{z}]_1 < \frac{1}{B^N} \frac{B^N - B^C}{\hat{B}^N - B^C}([\Delta^{-1}\mu]_2 - [\Delta^{-1}\mu]_1).$$

Here, $[\bar{z}]_j$ and $[\Delta^{-1}\mu]_j$ denote the jth entry of vectors $\bar{z}$ and $\Delta^{-1}\mu$, respectively.

The condition is immediately satisfied if the supply of the first security is at least as high as that of the second security ($[\bar{z}]_2 - [\bar{z}]_1 \leq 0$), as was the case in experiment 981007 (where $[\bar{z}]_2 = [\bar{z}]_1$), because the right-hand-side of the above inequality is always positive (among other things, $[\Delta^{-1}\mu]_2 > [\Delta^{-1}\mu]_1$, as mentioned before). The condition will be satisfied even if the supply of the first security is higher than that of the second security ($[\bar{z}]_2 - [\bar{z}]_1 > 0$; this is the situation in the other experiments) provided the harmonic mean risk aversion B^N is sufficiently low, which obviously is how shortsale constraints will be binding in the first place.

Table 4 reports the percentage of subjects for whom shortsale constraints on the riskfree security (Notes) were binding. In all experiments, this percentage is significant, lending credibility to the hypothesis that binding shortsale constraints explain the pattern in mean allocation prediction errors revealed in Table 3.[11]

Table 4 shows a pattern in the percentage shortsale-constrained subjects that is to be expected from rational subjects. Even if they are risk-neutral, subjects should take into account foregone opportunities to earn more if excessively risky positions force them to be barred from further trading (see the Appendix for an explanation of the bankruptcy rule). Hence, their actions ought to reveal a higher level of risk aversion in the early periods of an experiments. Table 4 confirms this: fewer subjects are shortsale-constrained in the first half of the experiment.

The above analysis does not apply to experiment 001130. Remember that three state securities were traded. Because Notes were also available, one security is redundant. We translated subjects' holdings of Notes and three state securities into the equivalent holdings in terms of Notes and state securities X and Y (the tables refer to them as securities A and B, respectively). Shortsale constraints in the original full-state-securities configuration, however, do not readily translate into the configuration in terms of Notes, A and B. For instance, holdings of the state security Z correspond to short positions in A and B and a long position in Notes. That is, a subject who is holding a long position in the state security Z may even be recorded as having a short position in one of the risky securities once the holdings are translated into holdings of A, B and Notes only. Not surprisingly, the estimated average allocation prediction errors display no clear pattern (see Table 4).

6 Conclusion

The development of a structural econometric test of a general equilibrium asset pricing model appropriate for experimental data constitutes the methodological contribution of this paper. The test is purposely different from those used for field data. It exploits cross-sectional size rather than time series length. Its stochastic structure is based

[11]Further corroborating evidence came from a set of experiments where uncertainty was eliminated, but quadratic utility was induced by means of nonlinearities in subjects' payoff functions. See [2] for a description. Where shortsale constraints were binding at equilibrium, analogous patterns in the estimated average allocation prediction errors emerged.

on errors of expected utility theory in predicting subjects' choices, instead of estimation error in the distribution of payoffs (which is known in an experimental context). Unlike field data tests, the ready availability of portfolio holdings allows ours to verify the predictions of general equilibrium theory in a comprehensive way: the link between prices and allocations is tested, instead of some necessary condition on prices alone.

While the shape of the empirical distribution of the core statistic of the structural test corresponds to its asymptotic distribution under the posited model, it is displaced, implying that there are more rejections than expected by chances. "Learning" cannot explain this, for there is no relationship between the frequency of rejections and time. Small-sample biases are evident in the positive correlation between the test statistic and the distance from CAPM pricing.

The structural estimation allowed us to discover patterns in the data that would be impossible to discern in a simple graphical or descriptive-statistic analysis. In particular, the structural estimation generated estimates of the mean error that expected utility theory makes in predicting allocations. A strong pattern emerged, indicating that our model was mis-specified in that it ignored short-sale constraints on the riskfree security. Such mis-specification obviously provides another reason for the high number of rejections.

In principle, one could extend the structural econometric tests to encompass a variation of our equilibrium model that allows for binding shortsale constraints. The ensuing nonlinearities, however, make such an extension nontrivial. Remember that each subject's risk tolerance had to be estimated as part of the structural test. Because only a few choices are observed per subject, risk tolerance parameters are estimated with significant error. We set up our test statistic, however, in such a way that errors in the estimation of risk tolerances average out across subjects, so that large-sample properties are unaffected. It is not clear how such a trick may be devised when there are nonlinearities induced by binding shortsale constraints.

In any event, before embarking on generalizations, one ought to ponder whether binding shortsale constraints are the sole source for the high frequency of rejections. We already mentioned another source: small-sample biases. But there is another complication. One would expect rejections to be related to the number of shortsale-constrained subjects. The more subjects are shortsale-constrained, the bigger the deviation from CAPM equilibrium. Figure 5, however, displays no pattern in the relation between the (logarithm of our) test statistic and the percentage subjects that are shortsale-constrained...

References

[1] Bossaerts, P. and C. Plott [2000], "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," Caltech working paper.

[2] Bossaerts, P., C. Plott and W. Zame [2002], "Prices And Portfolio Choices In Financial Markets: Theory and Experimental Evidence," Caltech working paper.

[3] Gallant, R. [1987], *Nonlinear Statistical Models*. New York: Wiley.

Appendix I: Asymptotic Behavior of A

Here follows a proof that A provides a consistent estimator of the asymptotic covariance matrix of $\sqrt{N}h_{N,t}$. Because their estimation errors have no effect asymptotically, we first re-write A as a function of the true risk tolerances (including their true harmonic mean):

$$\begin{aligned} A &= (B^N)^2 \Bigg\{ \mathrm{cov}(z_{n,t}) - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\ &\quad -\frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N)(z_{n,t} - \bar{z}_N)' \Bigg\}. \end{aligned}$$

Next, consider $\frac{1}{N}\sum_{n=1}^{N} E[z_{n,t}|p_t^N]$. Apply the definition of $\bar{z}_N$ – see (23) – to derive the following:

$$\begin{aligned} \bar{z}_N - \frac{1}{N}\sum_{n=1}^{N} E[z_{n,t}|p_t^N] &= \frac{1}{N}\sum_{n=1}^{N}(z_{n,t} - E[z_{n,t}|p_t^N]) \\ &= \frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t} \\ &\to 0, \end{aligned}$$

by the law of large numbers (which applies here, because of the cross-sectional independence of the errors $\epsilon_{n,t}$). Because the difference between $\frac{1}{N}\sum_{n=1}^{N} E[z_{n,t}|p_t^N]$ and $\bar{z}_N$ disappears asymptotically, we will substitute the latter for the former in expressions that depend on the former.

Invoking (22), compute the difference between A and the covariance matrix of $\sqrt{N}h_{N,t}$:

$$\begin{aligned} & A - NE[h_{N,t}\, h_{N,t}{}'\, |p_t^N] \\ &= (B^N)^2 \Bigg\{ \mathrm{cov}(z_{n,t}) - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\ &\quad -\frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N)(z_{n,t} - \bar{z}_N)' \Bigg\} \\ &\quad - (B^N)^2 \left\{ \frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\, \epsilon_{n,t}{}' \right\} \end{aligned}$$

$$
\begin{aligned}
= \quad & (B^N)^2 \Bigg\{ \mathrm{cov}(z_{n,t}) - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\
& -\frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N)(z_{n,t} - \bar{z}_N)' \Bigg\} \\
& - (B^N)^2 \Bigg\{ \frac{1}{N}\sum_{n=1}^{N}\left(z_{n,t} - \bar{z}_N\right)\left(z_{n,t} - \bar{z}_N\right)' \\
& - \frac{1}{N}\sum_{n=1}^{N}\left(z_{n,t} - \bar{z}_N\right)\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)' \\
& - \frac{1}{N}\sum_{n=1}^{N}\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)\left(z_{n,t} - \bar{z}_N\right)' \\
& + \frac{1}{N}\sum_{n=1}^{N}\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)' \Bigg\} \\
= \quad & (B^N)^2 \Bigg\{ \mathrm{cov}(z_{n,t}) - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\
& - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N)(z_{n,t} - \bar{z}_N)' \\
& - \mathrm{cov}(z_{n,t}) + \frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\
& + \frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N)(z_{n,t} - \bar{z}_N)' \\
& - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})(z_{n,t} - \bar{z}_N)(\mu - p_t^N)'\Delta^{-1} \\
& - \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}(\frac{1}{b_n} - \frac{1}{N}\sum_{\nu=1}^{N}\frac{1}{b_\nu})\Delta^{-1}(\mu - p_t^N))(z_{n,t} - \bar{z}_N)' \\
& + \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)' \\
& + \frac{1}{2}\frac{1}{N}\sum_{n=1}^{N}\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)\epsilon_{n,t}{}' \Bigg\} \\
= \quad & \frac{1}{2}(B^N)^2 \Bigg\{ \frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\left(E[z_{n,t}|p_t^N] - \bar{z}_N\right)'
\end{aligned}
$$

$$+\frac{1}{N}\sum_{n=1}^{N}\left(E[z_{n,t}|p_t^N]-\bar{z}_N\right)\epsilon_{n,t}{}'\Bigg\}.$$

Since $\frac{1}{N}\sum_{n=1}^{N}E[z_{n,t}|p_t^N]$ and $\bar{z}_N$ are indistinguishable asymptotically, let us replace $\bar{z}_N$ in the last expression. We obtain:

$$\begin{aligned} & A - NE[h_{N,t}\ h_{N,t}{}'\,|p_t^N] \\ = \quad & \frac{1}{2}(B^N)^2\left\{\frac{1}{N}\sum_{n=1}^{N}\epsilon_{n,t}\left(E[z_{n,t}|p_t^N]-\frac{1}{N}\sum_{\nu=1}^{N}E[z_{\nu,t}|p_t^N]\right)'+\frac{1}{N}\sum_{n=1}^{N}\left(E[z_{n,t}|p_t^N]-\frac{1}{N}\sum_{\nu=1}^{N}E[z_{\nu,t}|p_t^N]\right)\epsilon_{n,t}{}'\right\}.\end{aligned}$$

Once again, cross-sectional independence of the allocation prediction errors $\epsilon_{n,t}$ implies that the two terms in parentheses converge to zero, by the law of large numbers.

Appendix II: Description Of Experiments

Three assets could be traded. Two, labeled Securities A and B, had 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, was riskfree, and unlike the risky securities could be sold short, up to eight units. Note that markets are complete, which has particular implications for theoretical predictions. Among other things, this implies that, in equilibrium, shortsale of the risky securities has not restrictive. In contrast, subjects' ability to shortsell the riskfree security is potentially important even in equilibrium. Subjects were endowed with a certain number of securities A and B. The Notes were always in zero net supply.

Subjects could offer for sale or purchase, or trade, in each of the securities during periods of fifteen to twenty-five minutes, after which the state was announced and dividends were paid depending on the final holdings of the securities. The assets were then taken away, and subjects were given a fresh supply of the risky securities, as well as a certain amount of cash. The times between initial allocation of securities and cash and final realization of the dividends will be referred to as *periods*. Each experiment involved up to nine periods.

Cash was given at the beginning of every period. At the end of each period, subjects had to pay the experimenter for the initial cash as well as the initial allocation of risky securities. That is, initial cash and risky securities were given "on loan," and the end-of-period payment to the experimenter was therefore referred to as "loan repayment." This loan repayment created leverage, causing a magnification of the risk involved in the holding of securities A and B.

All accounting was 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 were also given an initial sign-up fee, theirs to keep if they succesfully finished the experiment (but fully exposed to any losses from holding risky securities).

Subjects were barred from further trading if they ran net cumulative losses more than two periods in a row. Such a bankruptcy rule induces risk aversion up to the penultimate period of an experiment even if subjects are risk neutral.

With the exception of experiment 001113, the payoff matrix was always the same, namely:

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

Note that the payoffs on securities A and B are negatively correlated. In other words, purchases of B can readily be used to diversify the risk of A.

Experiment 001113 had a different payoff matrix. While the number of states remained the same (3), we added a security (security C) and gave subjects a simple payoff matrix:

State	X	Y	Z
Security A	200	0	0
Security B	0	200	0
Security C	0	0	200
Notes	100	100	100

Note that there is redundancy: a portfolio of one unit of each of the risky securities generates the same payoff pattern as two Notes. The risky securities are of a type that is fundamental to general equilibrium asset pricing theory: they are state securities, paying a positive amount in one state, and nothing otherwise.

Trade took place over the internet, in a set of parallel, continuous, computerized open book markets, one for each security. The open book system worked as the ones used in field markets, such as the Paris Bourse, with one major exception: no orders could be hidden. The identity of the traders was not revealed; only an ID number identified orders in the book. The trading system, referred to as *Marketscape*, was developed at Caltech. It had been used succesfully in other experiments, also aimed at studying competitive equilibrium theory. The reader who would like to try out the mechanics of Marketscape should visit http://eeps3.caltech.edu/market-demo/.

The remaining data and parameters for the experiments are displayed in the table below. Notice the variation in initial endowments across experiments. With the exception of the first experiment, endowments also differed across subjects within an experiment. Subjects were not told other subjects' endowments (they were not even told how many subjects participated). This is a deliberate design feature. We wanted to avoid that subjects could easily deduce the nature of the aggregate risk (in the jargon of the most popular model: the payoff pattern of the market portfolio). Once aggregate risk is known, subjects could use standard models (like the CAPM) to price securities if they cared to do so. The purpose of the experiments is to generate the predictions of asset pricing theory, not through subjects' deliberate pricing on the basis of the theory that they may have heard about, but through the fundamental economic forces that the theory claims are behind its pricing results. Note also that subjects need not know the nature of aggregate risk for the predictions of general equilibrium theory in complete markets to be valid. In general equilibrium theory, prices are the only piece of marketwide information that agents are given.

States were drawn randomly. In the earlier experiments, states were drawn with replacement, meaning that the probability of occurence of each state remained 1/3 throughout. In the later experiments, states were drawn without replacement, as follows. Subjects were told to envision an urn with 18 balls, 6 of whom were marked X, 6 were marked Y, and 6 were marked Z. A ball is drawn at random. The letter marked on this ball determines the state, say X. Subjects were then told the new composition of the urn: 5 of X, 6 of Y and 6 of Z. The state for the subsequent period was determined through random drawing from this urn with 17 balls.

The table below summarizes the parameter settings for each experiment.

Date	Draw Type [a]	Subject Category (Number)	Signup Reward (franc)	Endowments A	B	C	Notes (franc)	Cash (franc)	Loan	Exchange Rate $/franc
981007	I	30	0	4	4	N/A[b]	0	400	1900	0.03
981116	I	23	0	5	4	N/A	0	400	2000	0.03
		21	0	2	7	N/A	0	400	2000	0.03
990211	I	8	0	5	4	N/A	0	400	2000	0.03
		11	0	2	7	N/A	0	400	2000	0.03
990407	I	22	175	9	1	N/A	0	400	2500	0.03
		22	175	1	9	N/A	0	400	2400	0.04
991110	I	33	175	5	4	N/A	0	400	2200	0.04
		30	175	2	8	N/A	0	400	2310	0.04
991111	I	22	175	5	4	N/A	0	400	2200	0.04
		23	175	2	8	N/A	0	400	2310	0.04
000804	I	8	175	5	4	N/A	0	400	2200	0.04
		7	175	2	8	N/A	0	400	2310	0.04
001113	I	29	125	7	13	9	0	400	2200	0.04
		30	125	8	11	12	0	400	2310	0.04
011114	D	21	125	5	4	N/A	0	400	2200	0.04
		12	125	2	8	N/A	0	400	2310	0.04
011126	D	18	125	5	4	N/A	0	400	2200	0.04
		18	125	2	8	N/A	0	400	2310	0.04
011205	D	17	125	5	4	N/A	0	400	2200	0.04
		17	125	2	8	N/A	0	400	2310	0.04

[a] I=states are drawn independently across periods; D=states are drawn without replacement, starting from an urn with six balls of each type (state).

[b] N/A = not applicable.

We did not have full control over the number of participants in the experiments. Subjects could 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 Marketscape website from which the experiment would be run, and execute some trades as practice. Subjects were given clear instructions, which they could read on the website for each experiment, and which we did not alter across experiments. These instructions included a description of some portfolio strategies. They did not mention a strategy whereby the subject is invited to gamble against the experimentor's random drawing of states.[12]

[12] The interested reader can visit http://eeps3.caltech.edu/market-011126 to view the instructions for market-011126, which include a

It should be pointed out that most subjects had familiarity with financial markets. Students were a mixture of U.S. and Canadian MBAs, economics undergraduates and undergraduates in the sciences and engineering. Students in 011126 had no prior experience with financial markets: they were mathematics majors of Sofia University in Bulgaria. Experiment 011126 was meant to reveal to what extent payment scale affects the results. The Bulgarian students were paid on the same dollar scale. Because their standard of living is less than 1/10th that of the U.S. and Canada, we effectively paid them more than 10 times as much as U.S. and Canadian students.

Bulgarian translation of the standard English instructions.

Table 1: Average Absolute Difference Between Market Sharpe Ratio and Maximal Sharpe Ratio (Average Based On Last Ten Transactions; Maximum In Parentheses).

Experiment	Period								
	1	2	3	4	5	6	7	8	9
981007	0.13	0.02	0.03	0.01	0.02	0.01			
	(0.71)	(0.31)	(0.34)	(0.27)	(0.13)	(0.05)			
981116	0.06	0.00	0.01	0.01	0.04	0.03			
	(2.74)	(0.16)	(0.02)	(0.09)	(0.16)	(2.04)			
990211	0.10	0.20	0.05	0.08	0.03	0.16	0.24		
	(13.17)	(0.67)	(0.29)	(0.39)	(0.09)	(0.21)	(0.56)		
990407	0.02	0.00	0.03	0.11	0.39	0.03	0.08	0.45	
	(0.35)	(0.37)	(0.17)	(0.34)	(0.55)	(0.53)	(0.22)	(0.67)	
991110	0.20	0.19	0.11	0.02	0.01	0.04			
	(0.74)	(0.30)	(0.31)	(0.23)	(0.22)	(0.11)			
991111	0.01	0.02	0.16	0.15	0.14	0.14	0.03	0.02	
	(0.49)	(0.35)	(0.20)	(0.30)	(0.27)	(0.29)	(0.33)	(0.13)	
000804	0.09	0.10	0.02	0.03	0.07	0.12	0.13		
	(0.33)	(0.20)	(0.26)	(0.05)	(0.10)	(0.14)	(0.40)		
001113	0.11	0.05	0.06	0.09	0.04	0.28	0.03	0.00	0.20
	(0.16)	(0.06)	(0.12)	(0.13)	(0.33)	(0.32)	(0.11)	(0.09)	(0.28)
011114	0.10	0.01	0.00	0.01	0.14	0.07	0.10	0.17	0.04
	(0.21)	(0.15)	(0.17)	(0.09)	(0.22)	(0.20)	(0.34)	(0.24)	(0.11)
011126	0.24	0.06	0.03	0.03	0.01	0.01	0.00	0.05	
	(4.28)	(0.53)	(0.14)	(0.14)	(0.03)	(0.29)	(0.36)	(0.07)	
011205	0.00	0.08	0.12	0.03	0.07	0.07	0.09	0.14	0.01
	(0.52)	(0.12)	(0.15)	(0.13)	(0.10)	(0.13)	(0.18)	(0.45)	(0.21)

Table 2: GMM Tests Of Equilibration

Experiment	Statistic	Periods 1	2	3	4	5	6	7	8	9
981007	χ_1^2	36.2	2.2	79.3	28.9	21.0	12.6			
	p-level	.00	.14	.00	.00	.00	.00			
	b^N (*10^{-3})	0.8	0.7	1.3	1.1	1.3	1.1			
	s.e. (*10^{-3})	0.0	0.1	0.0	0.1	0.0	0.0			
981116	χ_1^2	23.7	0.9	1.0	4.4	3.5	30.3			
	p-level	.00	.35	.32	.04	.06	.00			
	b^N (*10^{-3})	1.5	1.1	0.8	1.0	1.9	2.0			
	s.e. (*10^{-3})	0.1	0.1	0.1	0.1	0.1	0.1			
990211	χ_1^2	5.3	11.2	5.5	33.3	4.0	15.4	0.2		
	p-level	.02	.00	.02	.00	.04	.00	.69		
	b^N (*10^{-3})	1.1	1.5	1.4	2.8	1.2	1.5	-0.2		
	s.e. (*10^{-3})	0.1	0.1	0.1	*[a]	0.1	0.1	0.1		
990407	χ_1^2	7.5	0.7	13.8	116.7	†[b]	2.5	13.6	†	
	p-level	.01	.39	.00	.00	-	.62	.00	-	
	b^N (*10^{-3})	0.5	0.5	0.3	-0.3	2.8	0.3	0.2	0.9	
	s.e. (*10^{-3})	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.1	
991110	χ_1^2	197.5	72.8	31.6	7.4	2.7	6.9			
	p-level	.00	.00	.00	.01	.10	.01			
	b^N (*10^{-3})	3.0	2.4	1.9	1.2	1.1	1.4			
	s.e. (*10^{-3})	*	0.1	0.1	0.1	0.1	0.1			
991111	χ_1^2	4.8	1.5	114.4	61.4	36.3	43.4	31.6	30.8	
	p-level	.03	.23	.00	.00	.00	.00	.00	.00	
	b^N (*10^{-3})	0.4	0.3	1.7	1.4	1.2	1.0	1.3	1.3	
	s.e. (*10^{-3})	0.0	0.1	*	0.0	0.0	0.0	0.0	0.0	
000804	χ_1^2	31.3	20.4	9.6	7.5	25.3	12.1	4.9		
	p-level	.00	.00	.00	.01	.00	.00	.03		
	b^N (*10^{-3})	1.5	1.0	1.2	0.9	1.3	0.7	0.2		
	s.e. (*10^{-3})	0.0	0.0	0.1	0.1	0.0	0.1	0.1		
001113	χ_1^2	0.3	0.5	0.0	0.3	1.2	11.6	0.5	22.4	1.7
	p-level	.61	.49	.88	.60	.27	.00	.46	.00	.19
	b^N (*10^{-3})	0.1	-0.0	-0.0	0.4	0.3	0.7	0.3	0.6	0.4
	s.e. (*10^{-3})	0.2	0.2	0.2	0.2	0.3	0.1	0.1	0.1	0.1
011114	χ_1^2	37.5	2.3	2.2	5.4	17.6	15.4	10.7	21.8	4.4
	p-level	.00	.13	.14	.02	.00	.00	.00	.00	.04
	b^N (*10^{-3})	0.9	0.9	0.9	0.9	1.4	0.5	0.1	1.1	1.3
	s.e. (*10^{-3})	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.1	0.0
011126	χ_1^2	186.6	0.6	7.8	5.6	1.5	1.3	0.2	6.8	
	p-level	.00	.44	.01	.02	.23	.25	.65	.01	
	b^N (*10^{-3})	4.1	1.3	1.7	1.7	1.0	1.3	1.1	1.7	
	s.e. (*10^{-3})	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	
011205	χ_1^2	2.9	10.4	17.4	13.8	13.0	15.3	7.8	19.0	5.0
	p-level	.09	.00	.00	.00	.00	.00	.01	.00	.02
	b^N (*10^{-3})	0.8	1.8	1.7	1.5	1.9	1.3	0.6	0.6	0.7
	s.e. (*10^{-3})	0.0	0.1	0.1	0.1	0.1	0.1	0.0	0.0	0.0

[a] * denotes that the weighting matrix was not positive definite, and hence, standard errors could not be computed.

[b] † denotes negative χ^2 because weighting matrix was not positive definite.

Table 3: Average Estimated Allocation Prediction Errors Relative To Mean-Variance Optimal Demands

Experiment	Security	Periods								
		1	2	3	4	5	6	7	8	9
981007	A	2.03	1.29	2.23	2.20	1.66	1.19			
	B	-2.82	-1.23	-1.74	-2.59	-1.65	-1.22			
981116	A	1.77	0.55	0.65	1.03	0.51	1.43			
	B	-2.82	-1.23	-1.74	-2.59	-1.65	-1.22			
990211	A	1.66	2.55	1.89	1.60	1.77	2.40	-1.71		
	B	-2.01	-2.99	-2.27	-2.07	-1.54	-4.50	1.91		
990407	A	2.71	0.51	-3.59	15.25	4.29	-0.71	-6.52	4.11	
	B	-4.48	-0.95	6.23	-13.27	8.30	0.92	8.30	19.49	
981110	A	1.68	1.91	1.84	1.38	0.95	0.98			
	B	-3.65	-3.29	-2.82	-1.84	-1.30	-1.46			
991111	A	1.47	-1.84	2.30	2.18	2.59	2.86	1.51	1.11	
	B	-2.02	2.26	-2.58	-2.63	-2.65	-2.89	-1.85	-1.10	
000804[a]	A	2.16	2.30	1.15	1.42	2.02	3.23	6.11		
	B	-1.52	-2.05	-1.77	-1.54	-2.21	-3.39	-5.08		
001113	A	-5.37	-53.16	3.65	-1.64	4.12	-2.30	2.76	7.99	3.84
	B	-2.86	-23.10	0.76	-0.77	1.40	-0.44	1.79	4.22	1.96
011114	A	2.58	0.63	0.74	1.39	2.42	1.88	1.64	2.44	0.86
	B	-3.12	-0.86	-0.88	-1.27	-2.07	-1.82	-1.63	-2.67	-1.14
011126	A	2.27	-0.38	1.57	1.26	1.19	0.58	-0.42	1.48	
	B	-3.75	0.51	-2.10	-1.68	-1.53	-0.90	0.42	-1.66	
011205	A	0.64	1.60	2.37	1.27	1.72	2.40	1.95	3.00	1.04
	B	-0.75	-2.46	-2.76	-2.16	-2.60	-2.57	-2.49	-2.87	-1.18

[a] Security A pays 200 in state X and 0 otherwise; security B pays 200 in state Y and 0 otherwise.

Table 4: Percentage Shortsale-Constrained Subjects

Experiment	Period								
	1	2	3	4	5	6	7	8	9
981007	7	17	13	20	17	17			
981116	9	9	14	18	20	11			
990211	5	16	11	5	11	16	16		
990407	7	9	11	9	16	11	9	16	
991110	5	8	11	16	14	11			
991111	20	16	18	16	11	18	9	11	
000804	7	7	7	7	7	13	7		
001113	16	18	18	21	16	13	13	9	11
011114	5	17	14	21	23	12	14	17	19
011126	6	6	6	6	14	11	11	14	
011205	6	18	9	15	6	18	29	26	22

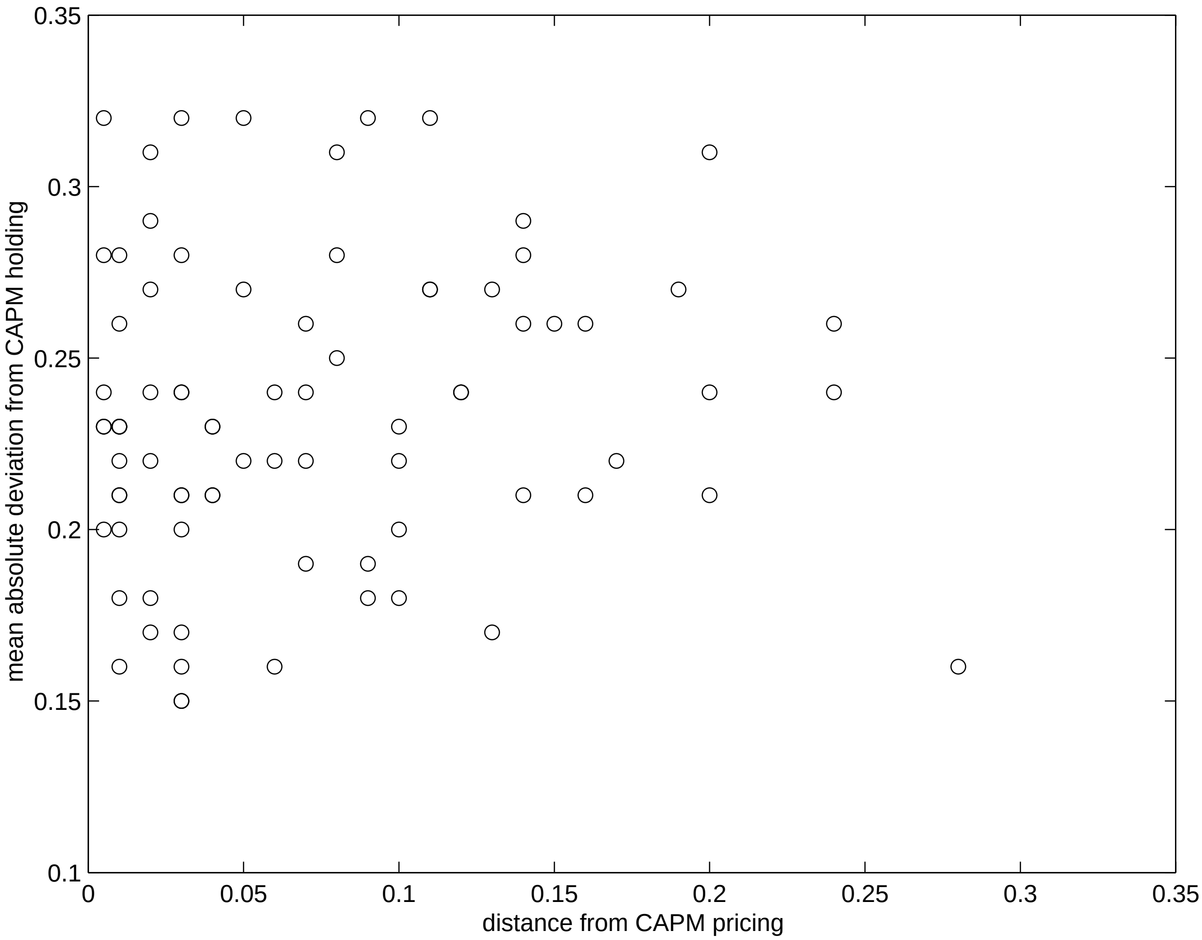

Figure 1: Plot of mean absolute deviations of subjects' end-of-period holdings from CAPM predictions against distances from CAPM pricing (absolute difference between market Sharpe ratio and maximal Sharpe ratio, based on last ten transactions), all periods in all experiments.

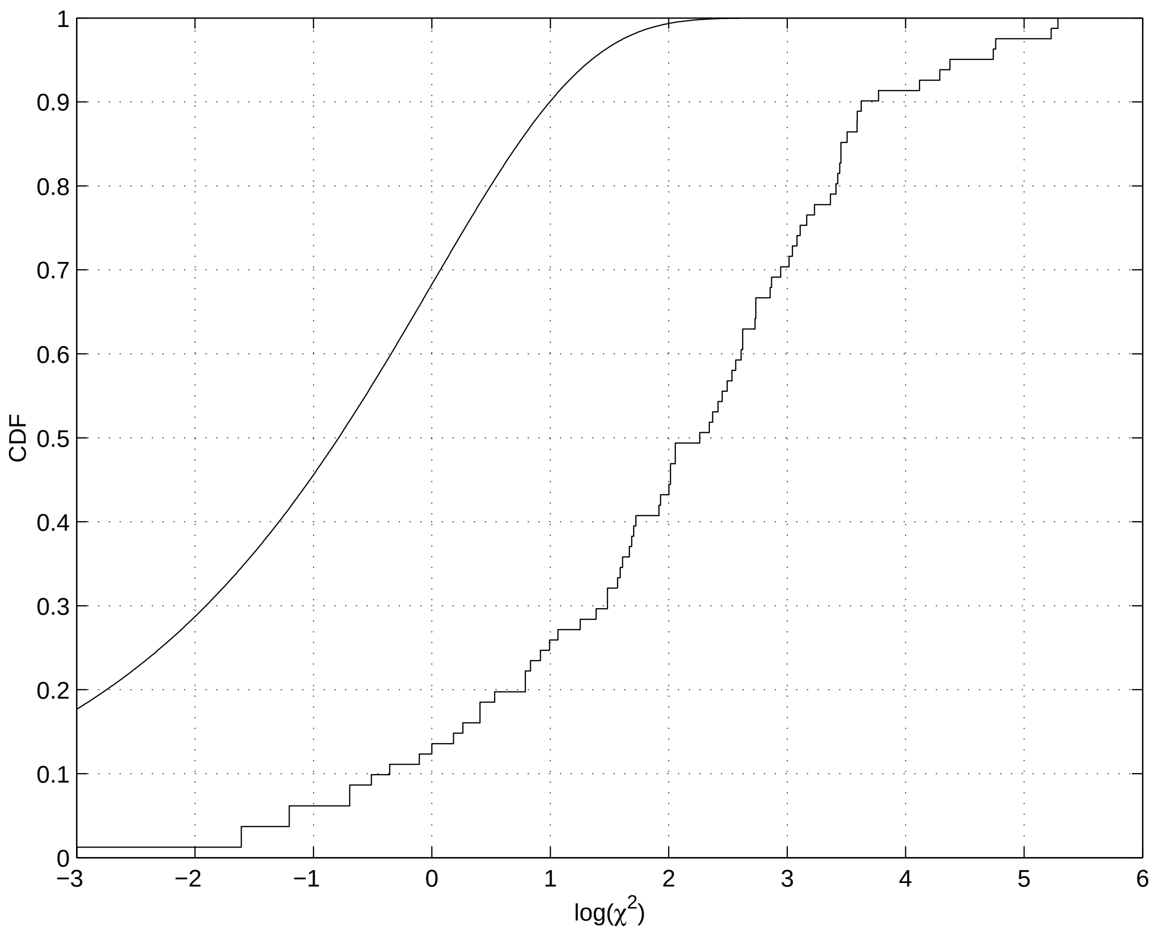

Figure 2: Empirical distribution of the GMM χ^2 statistics, all periods in all experiments (jagged line) against the theoretical asymptotic distribution under the hypothesis in (12).

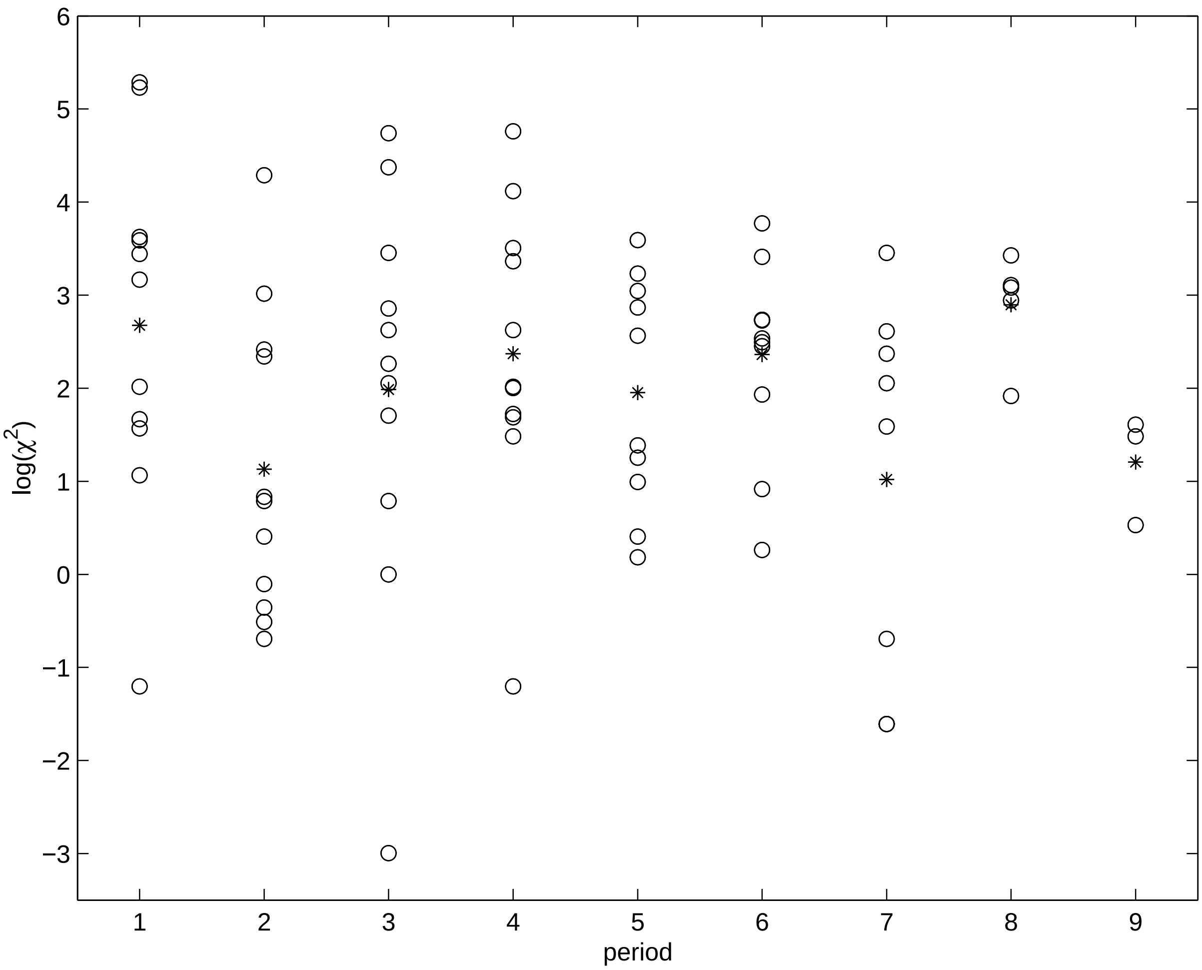

Figure 3: Plot of the logarithm of the GMM χ^2 statistics against period number. Stars represent period averages.

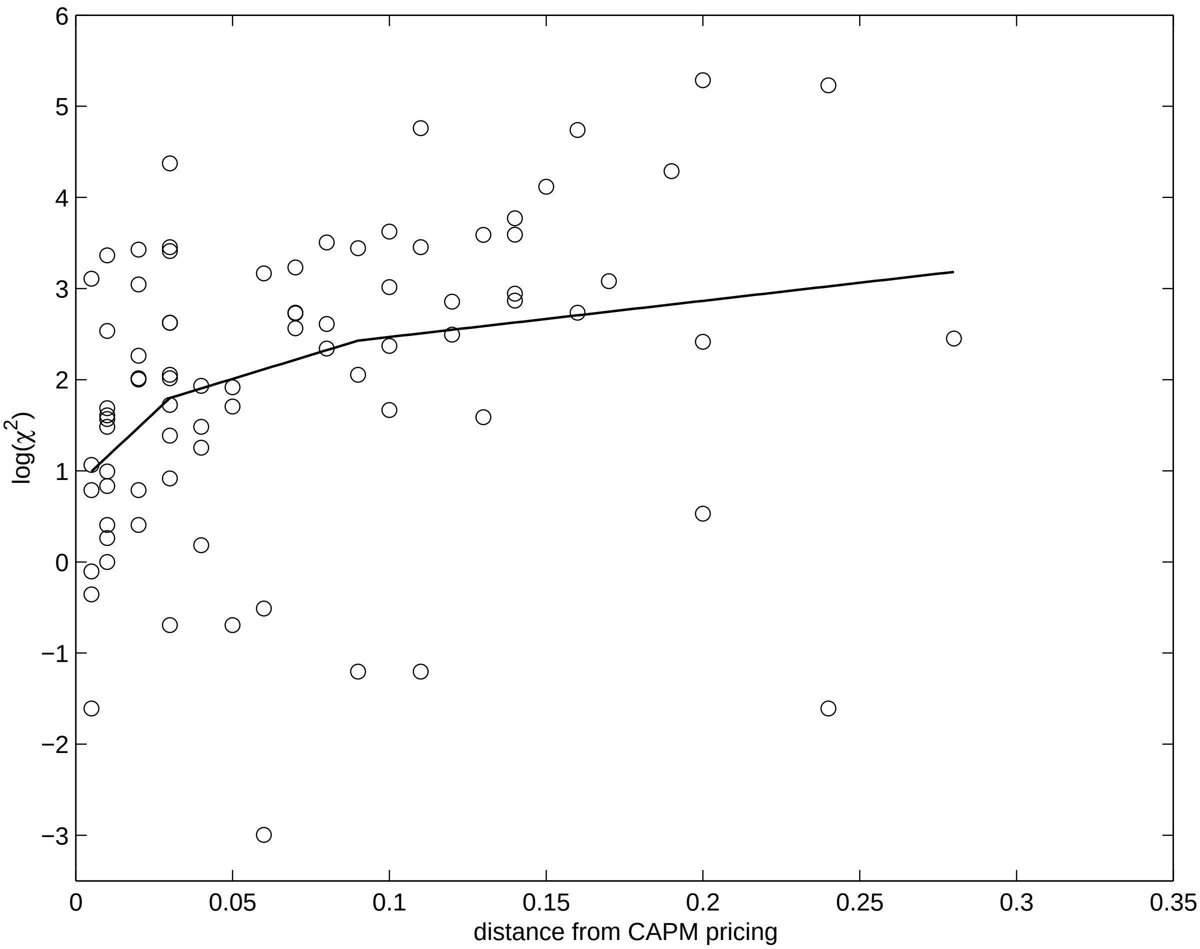

Figure 4: Plot of the logarithm of the GMM χ^2 statistics against distances from CAPM pricing (absolute difference between market Sharpe ratio and maximal Sharpe ratio, based on last ten transactions), all periods in all experiments. The solid line represents a piecewise linear spline fit. The scatter plot generates a (linear) correlation coefficient equal to 0.34 ($p < 0.01$).

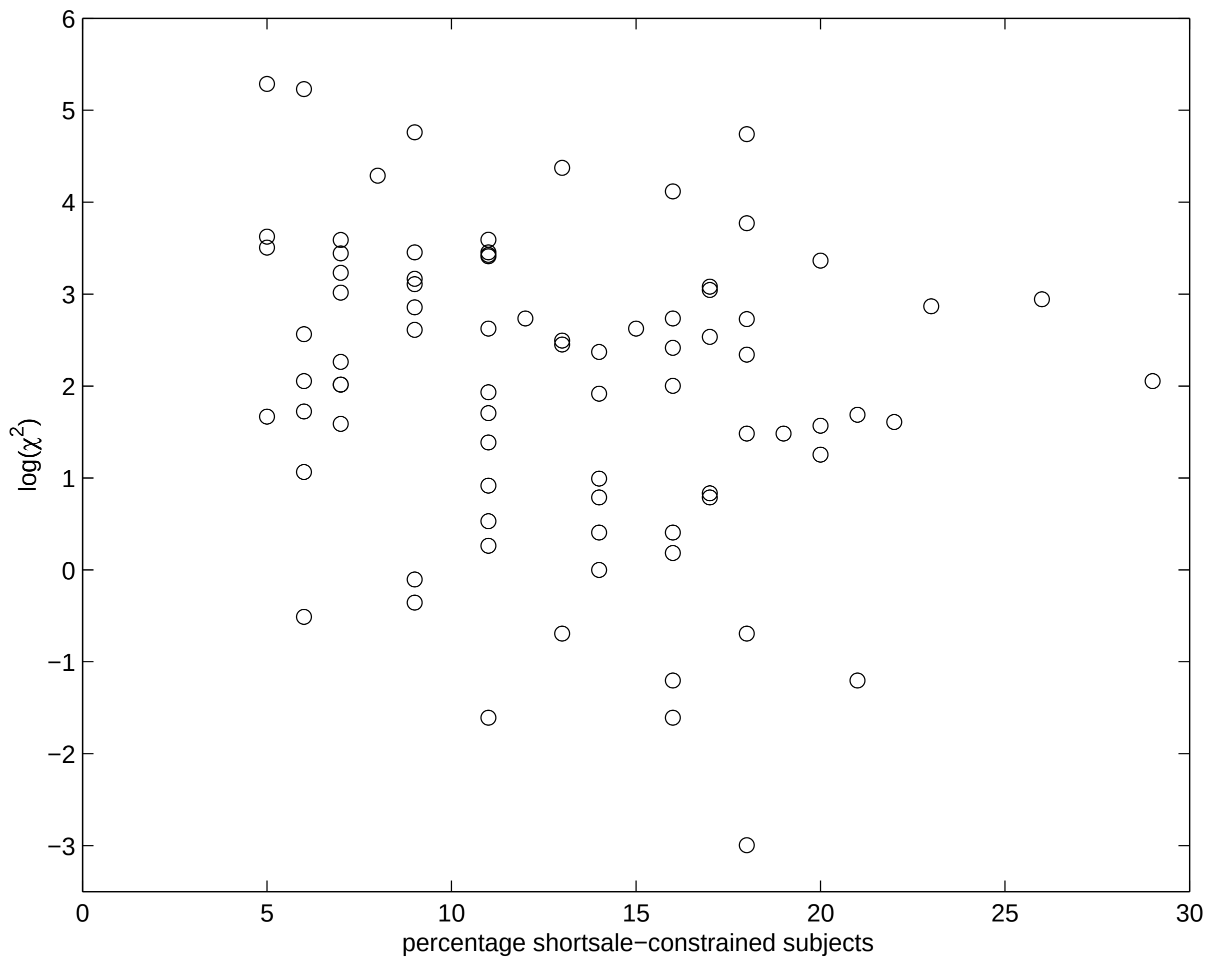

Figure 5: Plot of the logarithm of the GMM χ^2 statistic against the percentage of subjects for whom the constraint on shortsales of Notes is binding.