

Prices and Portfolio Choices
in Financial Markets:
Theory, Econometrics, Experiments*

Peter Bossaerts

California Institute of Technology
Centre for Economic Policy Research
Swiss Finance Institute

Charles Plott

California Institute of Technology

William R. Zame

UCLA
California Institute of Technology

July 1, 2003

This version March 18, 2007

*We thank seminar audiences at the Atlanta Finance Forum, Bachelier Finance Society International Congress, CMU, CEPR, CIDE, Erasmus University, Harvard University, Insead, IDEI, London Business School, NBER, RFS Conference on Behavioral and Experimental Finance, Rice University, SWET, UC Berkeley, UCLA, UC Riverside, University of Chicago GSB, University of Copenhagen, USC, University of Texas (Austin) for comments, and the R.G. Jenkins Family Fund, the National Science Foundation, the Caltech Laboratory for Experimental Economics and Political Science, the John Simon Guggenheim Foundation, the Social and Information Sciences Laboratory at Caltech, the UCLA Academic Senate Committee on Research and the Swiss Finance Institute for financial support. 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

Many tests of asset pricing models address only the pricing predictions — but these pricing predictions rest on portfolio choice predictions which seem obviously wrong. This paper suggests a new approach to asset pricing and portfolio choices, based on unobserved heterogeneity. This approach yields the standard pricing conclusions of classical models but is consistent with very different portfolio choices. Novel econometric tests link the price and portfolio predictions and take account of the general equilibrium effects of sample-size bias. The paper works through the approach in detail for the case of the classical CAPM, producing a model called CAPM+ ϵ . When these tests are applied to data generated by large-scale laboratory asset markets which reveal both prices and portfolio choices, CAPM+ ϵ is not rejected.

JEL Classification Numbers C91, C92, D51, G11, G12

Keywords experimental finance, experimental asset markets, risk aversion

1 Introduction

Many asset pricing models predict both asset prices and portfolio choices. For instance, the classical CAPM predicts portfolio separation (that all agents should hold the market portfolio), while Merton's (1973) continuous-time model and Connor's (1984) equilibrium version of APT both predict k -fund separation.¹ Forty years of econometric tests of such models present varying levels of support for the pricing predictions of these models — but even casual empiricism suggests that the portfolio choice predictions are badly wrong. Because the pricing predictions of these models are built on the portfolio choice predictions through on the assumption that portfolio choices are optimal given prices, the pricing predictions and the portfolio choice predictions seem inextricably linked, and it therefore seems hard to take the pricing predictions seriously while ignoring the portfolio choice predictions.

This paper suggests a new approach to asset pricing and portfolio choice that allows both to be taken seriously. Our approach maintains the essential parsimony of existing models, but adds unobserved heterogeneity. Applied to any of the models above, our approach would yield models identical pricing implications but be consistent with very different portfolio choices. (In particular, this approach suggests a rationale for testing pricing predictions without testing portfolio choice predictions.) Although our approach could be applied to a wide range of models, we choose here to apply it only to the classical CAPM, yielding a model we call CAPM $+\epsilon$. CAPM is a convenient and practical starting point because it is the simplest model that is consistent with observed prices in our experimental data; moreover, given the level of risk and reward in our experiments, other classical models would yield similar predictions.² CAPM $+\epsilon$ differs from the standard CAPM in assuming that demand functions of individual traders can be decomposed as sums of mean-variance components and idiosyncratic components, and that the idiosyncratic components are drawn from a distribution that has mean zero. These idiosyncratic components can be interpreted as reflecting unob-

¹Even the Fama-French three factor model (Davis, Fama and French, 2000) assumes that investors only hold combinations of three basic portfolios.

²See Judd and Guu (2001) for instance.

served heterogeneity of preferences; this is an approach similar to that used in much applied work. (The name CAPM+ ϵ is suggested by viewing true demand functions as perturbations of mean-variance demand functions.) We construct econometric tests of CAPM+ ϵ , and apply those tests to data generated by experimental financial markets in which both prices and portfolio choices can be observed. CAPM+ ϵ does well on this data: it is rejected (by either of our tests) at the 1% level in only one sample in eight, and at the 5% level in only two samples in eight.

We arrange our presentation in the following way. We begin (Section 2) by describing CAPM+ ϵ and providing a simple theoretical analysis to show that CAPM+ ϵ yields the standard pricing conclusions of classical models but is consistent with very different portfolio choices.³

With this simple theoretical analysis in hand, we next (Section 3) extend familiar analyses of models of individual choice to derive econometric tests of our model. We make use of the model assumption that individual demand functions can be decomposed into mean-variance components and idiosyncratic components and that the idiosyncratic components are random (which provides a source of variation) and from a distribution that has mean zero. Our tests are based on the generalized method of moments (GMM), adapted to large cross-sections, rather than long time-series. The tests are novel in a number of ways:

- Our tests link prices *and* choices.
- In the usual models of choice with unobserved heterogeneity, the null

³In a continuum economy, the idiosyncratic components of demand would *exactly* average out across the population and prices would be *exactly* as predicted by CAPM — but portfolio choices might be much different. In a large but finite economy, an appropriate form of the Law of Large Numbers implies that the idiosyncratic components of demand will (with high probability) *almost* average out across the population. Using this, we show that (with high probability) prices will be near to those predicted by CAPM — but again, portfolio choices may be much different. The argument is not difficult, but it is not trivial, because the competitive equilibrium correspondence is *not* continuous; hence a small perturbation might cause equilibria to disappear. The result we need depends on the fact that the CAPM equilibrium of the continuum economy is regular.

hypothesis is that the idiosyncratic components of demand have mean zero and are orthogonal to prices. In our setting, the idiosyncratic components of demand *functions* have mean zero and are orthogonal to prices. However, because demands influence equilibrium prices, the *realizations* of the idiosyncratic components of demand functions need *not* be orthogonal to prices. This induces a significant small-sample bias. Our tests accommodate this bias, which means that our null hypothesis reflects a *Pitman drift*. As a result, the asymptotic distribution of our GMM test statistic is *non-central* χ^2 . (Absent the small-sample bias, the asymptotic distribution would be *central* χ^2 .)

- To compute the weighting matrix for our GMM statistic, we need estimates of individual risk tolerances (inverses of risk aversion coefficients). Inspired by techniques introduced in McFadden (1989) and Pakes and Pollard (1989), we obtain individual risk tolerances using (unbiased) OLS estimation. Because the error averages out across subjects, this strategy enables us to ignore the (fairly large) error in estimating individual risk tolerances.

The next part of our presentation (Section 4) describes our experimental asset markets. Experimental asset markets are ideal for our purpose because they make it possible to observe (or control) many parameters that are difficult or impossible to observe in the historical data, including the market portfolio, the true distribution of returns, the information held by investors, and portfolio choices. In each of our experimental markets, 30 - 60 subjects trade riskless and 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 states are drawn independently at the end of each period. Subjects know the dividend structure (the payoff of each security in each state of nature) and the proba-

bility 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 (the aggregate supply of risky securities).

In Section 5, we present our analysis of the experimental data. We follow standard strategy, familiar from empirical studies of historical data, and use end-of-period prices and portfolio holdings, ignoring intra-period prices.⁴ Our experiments actually represent an environment that is closer to a static asset-pricing model than are typical studies of historical data, because our securities have only one-period lives. We can use liquidating dividends as security payoffs, while empirical tests of historical data usually take the relatively arbitrary end of the month as the end of the period, and use prices (including dividends collected during the month) as security payoffs.

Because the subject population is constant during an experiment, different periods within a single experiment do not represent independent draws. However, because subject populations of different experiments are disjoint, identical periods of different experiments do represent independent draws. We use this fact to construct multiple independent samples by using data *across* experiments. From these samples we construct empirical distributions of our test statistic, and use standard (Kolmogorov–Smirnov and Cramer–von Mises) tests to measure goodness-of-fit of these empirical distributions with a non-central χ^2 distribution function. The summary results are that the Kolmogorov–Smirnov and Cramer–von Mises goodness-of-fit tests reject CAPM+ ϵ at the 1% level in only one of eight samples and reject at the 5% level in only two of eight samples. These results are especially striking in comparison to other tests of asset-pricing models, which frequently reject at the 0.5% level. Because Monte Carlo methods are neither practical nor appropriate in this context, we provide an alternative power analysis based on perturbations of the existing data (rather than on artificial data, which Monte Carlo methods would utilize). This analysis suggests that our tests

⁴This is not to say that intra-period prices are of no interest; see Asparouhova, Bossaerts and Plott (2003) and Bossaerts and Plott (2004) for detailed discussions.

have significant power: relatively small perturbations of prices lead to many more rejections.

Following all of this, Section 6 concludes. Technical details are relegated to various appendices.

2 CAPM+ ϵ

In this Section we offer a model that makes pricing predictions close to that of the classical CAPM but is consistent with very different portfolio choices. Our starting point is suggested by the idea, familiar from applied work, that parametric specifications of preferences represent only a convenient approximation of the observed/expressed/true demand structure in the marketplace. We implement this idea by viewing observed demands as *perturbations* of “ideal” demands (hence the name CAPM+ ϵ). In principle, these perturbations might represent some combination of subject errors (in computing and implementing optimal choices), market frictions and unobserved heterogeneity. Because an adequate treatment of errors or market frictions would necessitate a fully stochastic model, which we are not prepared to offer, and because we have some evidence that errors and market frictions are not of most importance in our setting (but see Bossaerts, Plott and Zame (2002)), we focus here on unobserved heterogeneity.

Because the approach we follow is quite intuitive (there is only one small subtle issue), the following informal description is sufficient for our needs. To fix ideas, we first formalize equilibrium in asset markets, and recall the standard development of the classical CAPM in a particular context appropriate to our experiments. (Appendix A fills in some missing details.) We then present the formal model which adds perturbations to individual demands. (Appendix B presents a careful and rigorous justification.)

2.1 Equilibrium in Asset Markets

In our experiments, two risky assets and a riskless asset (Notes) are traded against Cash. Because Cash and Notes have the same payoffs, we treat them as redundant assets, and use a model with two risky assets and one riskless asset. Thus, in our model, investors trade *assets* A, B, N , which are claims to state-dependent consumption. In our experiments, there are 3 states of nature X, Y, Z . We write $\text{div } A$ for the state-dependent dividends of asset A , $\text{div } A(s)$ for dividends in state s , and so forth. If $\theta = (\theta_A, \theta_B, \theta_N) \in \mathbf{R}^3$ is a *portfolio* of assets, we write

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

for the state-dependent dividends on the portfolio θ .

There are I investors. Investor i is characterized by an endowment portfolio $\omega^i = (\omega_A^i, \omega_B^i; \omega_N^i) \in \mathbf{R}_+^2 \times \mathbf{R}$ of risky and riskless assets, and a strictly concave, strictly monotone utility function $U^i : \mathbf{R}^3 \rightarrow \mathbf{R}$ defined over state-dependent terminal consumptions. (To be consistent with our experimental design, we allow consumption to be negative.) Endowments and holdings of risky assets are constrained to be non-negative, but endowments and holdings of the riskless asset can be negative. In particular, risky assets cannot be sold short, but the riskless asset can be. Investors care about portfolio choices only through the consumption they yield, so given asset prices q , investor i chooses a portfolio θ^i to maximize $U^i(\text{div } \theta^i)$ subject to the budget constraint $q \cdot \theta^i \leq q \cdot \omega^i$.

An *equilibrium* consists of asset prices $q \in \mathbf{R}_{++}^3$ and portfolio choices $\theta^i \in \mathbf{R}_+^2 \times \mathbf{R}$ 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

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

- asset markets clear:

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

2.2 CAPM

The classical CAPM assumes investor i 's utility function for state-dependent wealth x has the (linear) mean-variance form:

$$U^i(x) = E(x) - \frac{b^i}{2} \text{var}(x) \quad (1)$$

where expectations and variances are computed with respect to the true probabilities, and b^i is absolute risk aversion.⁵ As usual, it is assumed 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).

The assumption of mean-variance utility, in conjunction with the other assumptions made above, are very nearly those of the *Capital Asset Pricing Model* (CAPM); the only difference is that we allow for short sales only of riskless assets, while the classical CAPM allows for short sales of both riskless and risky assets. As we show in Appendix A, however, if covariance of the risky assets A, B is negative — as it is in our experimental asset markets, and as we henceforth assume — short sales of risky assets are irrelevant, and the usual consequences of CAPM obtain. To describe these consequences briefly (see Appendix A for further detail), write $M = \sum \omega^i$ for the *market portfolio of all assets*, $m = \sum(\omega_A^i, \omega_B^i)$ for the *market portfolio of risky assets* and $\bar{M} = M/I$, $\bar{m} = m/I$ for the respective *per capita portfolios*. Write $\mu = (E(A), E(B))$ for the vector of expected dividends of risky assets and

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

⁵An frequently-used alternative assumption is that agents maximize expected utility with respect to a quadratic felicity function. At the scale of our experiments, the differences would be almost unobservable; we use the mean-variance specification only for econometric convenience. See Judd and Guu (2001).

for the covariance matrix of risky assets. It is convenient to normalize so that the price of the riskless asset is 1, so that $(p_A, p_B, 1) = (p, 1)$ is the vector of all asset prices. Abusing notation, write asset demands as functions of $(p, 1)$ or as functions of p , as is convenient. Write $Z^i(p)$ for investor i 's demand for all assets at prices p , and $z^i(p)$ for investor i 's demand for risky assets at prices p .

CAPM equilibrium prices \tilde{p} for risky assets and equilibrium demands are given by the formulas:

$$\tilde{p} = \mu - \left(\frac{1}{I} \sum_{i=1}^I \frac{1}{b^i} \right)^{-1} \Delta \bar{m} \quad (2)$$

$$z^i(\tilde{p}) = \frac{1}{b^i} \Delta^{-1} (\mu - \tilde{p}) \quad (3)$$

(The quantity $\left(\frac{1}{I} \sum \frac{1}{b^i} \right)^{-1}$ is frequently called the *market risk aversion* and $\left(\frac{1}{I} \sum \frac{1}{b^i} \right)$ is frequently called the *market risk tolerance*.) Because the demand and pricing formulas involve individual risk aversions, which are not directly observable, they are not testable. However, the following immediate consequences of these formulas *are* testable.

- **Mean-Variance Efficiency** The market portfolio m of risky assets is mean-variance efficient; that is, the expected excess return $E(\text{div } m) - q \cdot m$ on the portfolio m is highest among all portfolios having variance no greater than $\text{var}(\text{div } m)$.^{6,7}
- **Portfolio Separation** All investors hold a portfolio of risky assets that is a non-negative multiple of the market portfolio m of risky assets.

All of the predictions derived above depend on the assumption that investors are strictly risk averse: $b^i > 0$. It is not obvious that subjects will

⁶Because M, m differ only by riskless assets, the entire market portfolio M is also mean-variance efficient.

⁷There are other pricing implications of CAPM (for instance, linear beta pricing rule) but, as shown by Roll (1977), these are all consequences of Mean-Variance Efficiency.

display strict risk aversion in a laboratory setting. However, this is ultimately an empirical question, not a theoretical one. As we shall see later, our own data suggest strongly that individuals are risk averse. This is not a new finding; see Holt and Laury (2002) for instance.

2.3 Perturbations: CAPM+ ϵ

The classical CAPM assumes that investor i 's utility function U^i has the mean-variance form (1). However, we can always write an arbitrary utility function U^i as a perturbation of a mean-variance utility function \tilde{U}^i , and hence can view investor i 's observed/expressed/true demand function z^i as a perturbation of a mean-variance demand function \tilde{z}^i :

$$z^i(p) = \tilde{z}^i(p) + \epsilon^i(p)$$

Hence, we can model the observed economy as a perturbation of an “ideal” economy which differs only in that investors have mean-variance utilities. (Note that in the ideal economy, investor preferences differ only by the coefficient of absolute risk aversion, but that in the observed economy, investor preferences may be arbitrary.) Write D, \tilde{D} for the mean market excess demand functions for risky assets in the observed economy and in the ideal economy. By definition:

$$\begin{aligned} D(p) &= \frac{1}{I} \sum (z^i(p) - \omega^i) \\ &= \frac{1}{I} \sum (\tilde{z}^i(p) + \epsilon^i(p) - \omega^i) \\ &= \frac{1}{I} \left(\sum \tilde{z}^i(p) - \sum \omega^i \right) + \frac{1}{I} \sum \epsilon^i(p) \\ &= \tilde{D}(p) + \frac{1}{I} \sum \epsilon^i(p) \end{aligned}$$

All this is simply formal manipulation. The economic content of our model is in the following two assumptions:

- i) The characteristics (asset endowments ω^i and demand functions z^i) of investors in the observed economy are drawn independently from some distribution of characteristics.

- ii) The perturbations ϵ^i are drawn independently from a distribution with mean zero.

The first of these assumptions is innocuous; the second has real bite. To see the implications of these assumptions, note first that, by definition, an equilibrium price is a zero of mean market excess demand. Thus, if the mean perturbation $\frac{1}{I} \sum \epsilon^i$ is identically zero then equilibrium prices in the observed economy and in the ideal economy will exactly coincide. More generally, if the mean perturbation $\frac{1}{I} \sum \epsilon^i$ is uniformly small, then equilibrium prices in the observed economy and in the ideal economy will nearly coincide. (Because the equilibrium correspondence is not lower hemi-continuous in the parameters of an economy, this is not entirely obvious; the proof relies on the fact that the CAPM equilibrium is regular; see Appendix B.) Because the perturbations are drawn independently from a distribution with mean zero, a suitable version of the Strong Law of Large Numbers in an appropriate function space will guarantee that if the number I of investors is sufficiently large then, with high probability, the mean perturbation $\frac{1}{I} \sum \epsilon^i$ will be uniformly small. In view of CAPM, the market portfolio of the ideal economy will be mean-variance efficient at the equilibrium price \tilde{p} for the ideal economy. Because the market portfolio for the observed economy is the same as the market portfolio for the ideal economy, if the number I of investors is large then, with high probability, the *market portfolio for the observed economy will be approximately mean-variance efficient at the equilibrium price for the observed economy*. Of course, individual portfolio choices in the observed economy need bear no obvious relationship to individual portfolio choices in the ideal economy; in particular, because the perturbations ϵ^i need not be small, approximate portfolio separation need not hold in the observed economy.

Perhaps the most important feature of this model is that it provides a mechanism that leads to mean-variance efficiency of the market portfolio even though *no single investor* chooses a mean-variance optimal portfolio. (In the classical CAPM of course, the market portfolio is mean-variance optimal because *every* investor chooses a mean-variance optimal portfolio.)

Because the pricing conclusions in our model are driven by the Strong

Law of Large Numbers, our model suggests that the likelihood that CAPM pricing will be observed is increasing in the number of market participants. Experimental evidence for this suggestion can be found in Bossaerts and Plott (2002).

3 Structural Econometric Tests

In this Section, we construct a structural econometric test of CAPM+ ϵ . Our approach has a number of novel features, some of which distinguish it from familiar econometric approaches to historical price data:

- The usual approaches rely entirely on market prices; our approach links market prices and individual holdings.
- In the usual approaches, the source of randomness is the error in estimation of the distribution of returns. In our approach, the source of randomness is the deviations of observed choices from ideal mean-variance-optimal choices.
- In the usual approach, GMM (Generalized Method of Moments) is used to construct an estimator that has good properties for long time series. In our approach GMM is adapted to construct an estimator that has good properties for large cross sections.⁸
- The usual approach is to test a theory on each sample separately, and then aggregate the results. Our approach is to construct samples across experiments, use our theory to infer the class of distribution to which our test statistic should belong, and use measures of goodness-of-fit to determine whether the empirical distribution of our test statistic on these samples is generated by a member of this class.

⁸We use this approach for two reasons. The first is that we do not have long time series — our experiments are only 6-9 periods long — but we do have a large cross-section — each experiment involves 30-60 subjects. The second is that the periods of our experiment do not represent independent draws, because they are populated by the same subjects.

Our approach requires estimates of risk aversion for each individual. To obtain such estimates, we use observed choices in *other* periods of a given experiment. Because the number of periods is small, such estimates are necessarily inaccurate; however, our estimation procedure is such that errors across subjects tend to cancel out.⁹

Our approach is inspired by econometric analysis of choice in panel data, but there is an important difference. In traditional analysis of models of individual choice, the error term is assumed to be independent of the explanatory variables. In CAPM+ ϵ , the error terms are perturbations of demand functions, and the explanatory variables include prices — but the definition of equilibrium entails that individual demands sum to supply, so that the perturbations *cannot* be independent of prices. This is a consequence of the fact that we have a *finite* population. In an infinite population, independent draws of perturbations from a distribution having mean zero would yield an aggregate perturbation that is identically zero, hence independent of prices. We therefore introduce a *Pitman drift* to capture the finite-sample bias in the large-sample analysis of our test statistic, which leads to an asymptotic distribution of our GMM statistic that is *non-central* χ^2 . Because we do not have detailed information about the perturbation terms, we do not know the value of the non-centrality parameter. To get around this problem, we will exploit the fact that our various experiments provide several independent replications with the same market parameters but non-overlapping populations, and thus represent independent draws. Each of these replications generates a sample. We can then test whether the empirical distribution of the GMM test statistics across samples is generated by some member of the family of non-central χ^2 distributions; this will be our test of the CAPM+ ϵ .

⁹This approach is reminiscent of the method used to obtain consistent standard errors in method of simulated moments with only a limited number of simulations per observation; see McFadden (1989) and Pakes and Pollard (1989).

3.1 The Null Hypothesis

We focus for the moment on an economy \mathcal{E}_t representing a single period t of a single experiment, in which there are I subjects/investors. Investor i is characterized by an endowment ω^i and a demand function for risky assets z_t^i .¹⁰ As in Section 2, we write μ for the vector of mean payoffs of the risky assets, Δ for the covariance matrix of payoffs of risky assets, and \bar{m}^I for the per capita market portfolio of risky securities. (We write z_t^i with a time subscript and \bar{m}^I with a population superscript as reminders that we focus on a particular period of a particular experiment. However, we do not subscript or superscript the parameters ω^i , μ , Δ because they do not depend on the particular period or experiment.) Because endowments ω^i are fixed throughout the experiment, we suppress them in what follows.

For each i , let b_t^i be the coefficient of risk aversion that most closely matches investor i 's end-of-period asset choices in *other periods* of the same experiment, and let \tilde{z}_t^i be the demand function for risky assets for an ideal investor having the same endowment as trader i and a mean-variance utility function as in (1) with coefficient of risk aversion b_t^i . The difference between the observed demand function z_t^i and the ideal demand function \tilde{z}_t^i is the *perturbation* or *error*:

$$\epsilon_t^i = z_t^i - \tilde{z}_t^i \tag{4}$$

Let $\tilde{\mathcal{E}}_t$ be the ideal economy populated by these mean-variance traders. Write

$$B_t^I = \left(\frac{1}{I} \sum_{i=1}^I \frac{1}{b_t^i} \right)^{-1},$$

for the market risk aversion for the ideal economy $\tilde{\mathcal{E}}_t$. (We use the superscript I to emphasize that we have an economy with I investors.)

The assumptions of CAPM hold in the ideal economy $\tilde{\mathcal{E}}_t$, so write \tilde{p}_t^I for the CAPM equilibrium price (again emphasizing dependence on I , t). At equilibrium, per capita demand for risky assets must equal the per capita

¹⁰As in Section 4 and Appendix B, we view the characteristics of the subjects as independent draws from a population with a given distribution.

market portfolio of risky assets, so rewriting equation (3) in the present notation yields

$$\tilde{z}_t^i(\tilde{p}_t^I) = \frac{1}{b_t^i} \Delta^{-1}(\mu - \tilde{p}_t^I) \quad (5)$$

As we show in Appendix A, it follows that $\tilde{p}_t^I = \mu - B_t^I \Delta \bar{m}^I$, and (assuming that $\|p_t^I - \tilde{p}_t^I\|$ is not too large):

$$\tilde{z}_t^i(p_t^I) = \frac{1}{b_t^i} \Delta^{-1}(\mu - p_t^I) \quad (6)$$

Assuming that end-of-period prices p_t^I are actually equilibrium prices for the economy \mathcal{E}_t , per capita demand must equal the per capita market portfolio:

$$\frac{1}{I} \sum z_t^i(p_t^I) = \bar{m}^I \quad (7)$$

Summing (5) and (6) over all investors i , combining with (7) and doing a little algebra yields the following relationship:

$$p_t^I = \tilde{p}_t^I + B_t^I \Delta \frac{1}{I} \sum_{i=1}^I \epsilon_t^i(p_t^I) \quad (8)$$

Note that prices p_t^I appear on *both* sides of this equation so it is *not* a formula for equilibrium prices.

A familiar intuition from applied work would suggest as null hypothesis the following statement:

The realizations $\epsilon_t^i(p_t^I)$ (of the perturbations ϵ_t^i at equilibrium prices p_t^I) are mutually independent across i , given p_t^I , and

$$E[\epsilon_t^i(p_t^I) | p_t^I] = 0 \quad (9)$$

(Keep in mind that the perturbations ϵ_t^i are *functions*, but that only the realizations $\epsilon_t^i(p)$ at some prices can possibly be observed.) This null hypothesis would lend itself readily to testing by means of the *Generalized Method of Moments statistic* (GMM, or minimum χ^2 statistic). In our setting, however, this would be the *wrong* null hypothesis: if the number of investors is

finite, market clearing implies that the perturbation terms *cannot* be independent of prices. Hence $E[\epsilon_t^i | p_t^I]$, the mean of the perturbations conditional on prices, may be different from zero even though $E[\epsilon_t^i] = 0$, the unconditional mean of the perturbations, is zero. Of course, $E[\epsilon_t^i | p_t^I] \rightarrow 0$ as $I \rightarrow \infty$ — perturbations have asymptotical conditional mean 0 — but we have only a finite sample, and we must take that into account.

Instead, we take as null hypothesis the following:

The conditional means of perturbations exhibit *Pitman drift*: for some λ ,

$$\lim_{I \rightarrow \infty} \sqrt{I} \left(E[\epsilon_t^i | p_t^I] \right) = \lambda \quad (10)$$

Here we view the economy as a draw of I investors from a fixed distribution of investor characteristics. So the expectation is taken over all investors in a particular draw of investors, conditional on equilibrium prices for that particular draw of investors, and then over all draws of investors; finally, we take the limit as the size of the draw tends to infinity. Under Pitman drift, the asymptotic distribution of the usual GMM statistic is non-central χ^2 with non-centrality parameter λ^2 . Because λ is unknown, CAPM+ ϵ *cannot* be tested on a single sample (a single period) but it *can* be tested on the behavior of the GMM statistic *across* samples, because the *form of its distribution* — non-central χ^2 — is known.

3.2 Specifics of the GMM Statistic

Define

$$h_t^I(\beta) = \beta \frac{1}{I} \sum_{i=1}^I z_t^i(p_t^I) - \Delta^{-1}(\mu - p_t^I). \quad (11)$$

We continue to use the superscript I to make explicit the dependence on the size of the drawn economy. Keep in mind that $h_t^I(\beta)$ depends on the particular draw, and therefore is a random variable. Now let β^I be the solution to the minimization problem:

$$\min_{\beta} [\sqrt{I} h_t^I(\beta)^T] W^{-1} [\sqrt{I} h_t^I(\beta)], \quad (12)$$

where W is a symmetric, positive definite weighting matrix (to be chosen below). The dependence of the solution on I is made explicit because we are interested in its asymptotic distributional characteristics as $I \rightarrow \infty$.

Under our null hypothesis, $h_t^I(\beta)$ is asymptotically zero in expectation when $\beta = B_t^I$. (To see this, note that:

$$\begin{aligned} h_t^I(B_t^I) &= B_t^I \frac{1}{I} \sum_{i=1}^I z_t^i(p_t^I) - \Delta^{-1}(\mu - p_t^I) \\ &= B_t^I \frac{1}{I} \sum_{i=1}^I \left[z_t^i(p_t^I) - \frac{1}{b_t^i} \Delta^{-1}(\mu - p_t^I) \right] \\ &= B_t^I \frac{1}{I} \sum_{i=1}^I \epsilon_t^i(p_t^I) \end{aligned} \tag{13}$$

Hence, $E[h_t^I(B_t^I)|p_t^I] = B_t^I \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i|p_t^I] \rightarrow 0$, as asserted.) The solution of (12) therefore defines a *GMM estimator of the market risk aversion*: it generates the value β which makes the sample version of the expectation in (13) as close as possible to zero, the (asymptotic) theoretical value of this expectation when $\beta = B_t^I$.

Because there are two risky assets, random variation in finite samples ensures that at β^I the distance from zero of the sample version of (13) is almost surely strictly positive. Our criterion function defined in equation (12) will be strictly positive in large samples as well, because the sample version of (13) is scaled by the factor \sqrt{I} . It has a well-defined asymptotic distribution. With the right choice of weighting matrix W , at its optimum β^I , our criterion function will be χ^2 distributed with one degree of freedom (the number of risky assets minus one) and with non-centrality parameter λ^2 . Hence, our criterion function defines a *GMM test of goodness-of-fit*.

3.3 Economic Interpretation of the GMM Test

It is illuminating to interpret the minimization that is part of the GMM test in terms of portfolio optimization. Because the weighting matrix W is required to be symmetric and positive definite, our GMM test verifies whether

the vector in (11) is zero. (To see this, note that $\frac{1}{I} \sum_{i=1}^I z_t^i$ in (11) is the mean demand for risky securities; at equilibrium prices, this must equal the mean market portfolio. Hence, if $h_t^I = 0$, the first-order conditions for mean-variance optimality are satisfied.) Thus the market portfolio will be optimal for an agent with mean-variance preferences and risk aversion parameter β , so *our GMM test verifies mean-variance optimality of the market portfolio*. Of course, verifying mean-variance optimality of the market portfolio is the usual way of testing CAPM on field data. In the usual field tests, however, distance from mean-variance efficiency is measured as a function of the error in the estimation of the distribution of payoffs; here we measure distance as a function of the weighting matrix W .

We define the weighting matrix W to be the asymptotic covariance matrix of

$$\sqrt{I} [h_t^I(B_t^I)] = B_t \sqrt{I} \left[\frac{1}{I} \sum_{i=1}^I \epsilon_t^i \right] \quad (14)$$

(see equation (13)). W is proportional to the asymptotic covariance matrix of the perturbations, so *our GMM statistic measures distance from CAPM pricing in terms of variances and covariances of the perturbations*. Allocational dispersion is the source of errors, not randomness in the estimation of return distributions. Our test thereby links prices to individual allocations, and thus provides a more comprehensive test of equilibrium than field tests – which rely only on prices or returns.

3.4 Estimating the Weighting Matrix W

For the necessary asymptotic distributional properties to obtain, the weighting matrix W should be estimated from the sample covariance matrix of the perturbations across subjects. Perturbations depend on individual risk tolerances $1/b_t^i$. Using an asymptotically (as $I \rightarrow \infty$) unbiased estimator, we obtain individual risk tolerances from portfolio choices across all periods in an experimental session except the period t on which the GMM test is performed. From the estimated risk tolerances, we compute individual perturbations for period t and, from those, we estimate W .

Since the number of periods in an experimental session (T) is small, the error in estimating risk tolerances may be large. However, because we use an asymptotically unbiased estimator of risk tolerances, the Law of Large Numbers implies that population means of the estimated risk tolerances converge to true population means. Moreover, since risk tolerance in period t is estimated from observations in periods other than t , the error in estimating an individual risk tolerance and that individual's perturbation for period t will be orthogonal, provided individual perturbations are independent over time. We write our estimator of W in such a way that we can exploit these two properties and ensure consistency even for fixed T . Appendix C discusses our procedure in more detail.

4 Experiments and Experimental Data

In this Section, we describe the design of our experiments and the data they generate. As will be seen, the experimental data illustrate the price-allocation paradox: there is no correlation between the conformity of prices with predictions and the conformity of choices with predictions. In particular: in some periods the market portfolio is mean-variance efficient — but choices never display portfolio separation.

4.1 Experimental Design

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

Each experimental session of approximately 2-3 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 riskless and risky assets and Cash. The endowments of risky assets and Cash are non-negative. Subjects are also

given loans, which must be repaid at the end of the period; we account for these loans as negative endowments of Notes. 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). 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. 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. (In some experiments, subjects were also given a bonus upon completion of the experiment.) 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. In effect, therefore, consumption in a given period can be negative. (The bankruptcy rule was seldom triggered.)

Subjects know their own endowments, and are informed about asset payoffs in each of 3 states of nature X, Y, Z , and of the objective probability distribution over states of nature. Randomization was achieved by drawing balls from an urn or by the use of a random number generator. In some experiments, states of nature for each period were drawn independently from the uniform distribution. In the remaining experiments, states were not drawn independently: either balls were drawn without replacement from an urn initially containing 6 balls for each state, or states were drawn by a random number generator according to the corresponding distribution. In each treatment, subjects were informed as to the procedure. Subjects were *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 were 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; indeed, since payoffs and probabilities are explicitly known, information provided to subjects is perhaps more than in these or other stock

markets.¹¹ We did not provide information about the market portfolio, so that subjects could not easily deduce the nature of aggregate risk.¹² Recall 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 — hence optimize with respect to — *only* payoffs, probabilities, market prices and their own preferences and endowments.

In the experiments reported here, there were three states of nature X, Y, Z . 1 unit of Cash is 1 franc in each state of nature; the state-dependent payoffs of assets (in francs) are recorded in Table 1 below.

Table 1: Asset Payoffs

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

The remaining parameters for the various experiments are displayed in Table 2. Each experiment is identified by the year-month-day on which it was conducted. Note that the social endowment (the market portfolio) and the distribution of endowments differ across experiments. Since equilibrium prices and choices depend on the social endowment (the market portfolio) and on the distribution of endowments, as well as on the preferences of investors, there is every reason to expect equilibrium prices to differ across experiments. Indeed, because subject preferences may not be constant across periods (because of changing probabilities, or wealth effects, or bankruptcy, or fear of bankruptcy), there is every reason to expect equilibrium prices to differ across periods in a given experiment. Note that, given the true

¹¹Recall that CAPM — and many other asset pricing models — assume that investors know the true distribution of returns.

¹²If the nature of aggregate risk were known, subjects might have used a standard model — CAPM, for instance — to *predict* prices, rather than taken observed prices as given. Hence CAPM pricing might have occurred because subjects *expected* it to occur.

probabilities, $\text{cov}(A, B) < 0$; as we shall see later, this simplifies the theory.

Subjects were given clear instructions, which included descriptions of some portfolio strategies (but no suggestions as to which strategies to choose).¹³ Most of the subjects in these experiments had some knowledge about economics in general and about 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 the experiment 011126, for which the subjects were undergraduates at the University of Sofia (Bulgaria), subjects may have been less knowledgeable. No subjects participated in more than one experimental session, so the subject populations have no overlap.

4.2 Prices and Choices

In our experiments, we observe and record every transaction. However, we focus here only on the ends of periods: that is, on the prices of the last transaction in each period and on individual holdings at the end of each period.^{14,15} Our focus on end-of-period prices and holdings is parallel to that of most empirical studies of historical data, which typically consider only beginning-of-month and end-of-month prices, and ignore prices at all intermediate dates. (Of course, the historical record provides little information about holdings.) In historical data, there is uncertainty at the beginning of each month about what prices — used as proxies for payoffs — will be at the end of each month. In our experiments, there is uncertainty at the end of each period about what state will be drawn and hence about what payoffs will be. (It is important to keep in mind that, although trading in

¹³Complete instructions and other details are available at <http://eeps3.caltech.edu/market-011126>; use anonymous login: ID 1, password a.

¹⁴The end of the period is in some ways a bit arbitrary; other possibilities might have been equally sensible. For example, we might have chosen instead to focus on averages over the last 10 seconds of each period.

¹⁵A complete record of every transaction in every experiment is available at <http://www.hss.caltech.edu/~pbs/BPZdata>.

Table 2: Experimental Parameters

Date	Draw Type ^a	Subject Category (Number)	Bonus Reward (franc)	Endowments			Cash (franc)	Exchange Rate \$/franc
				A	B	Notes ^b		
981007	I	30	0	4	4	-19.0	400	0.03
981116	I	23	0	5	4	-20.0	400	0.03
		21	0	2	7	-20.0	400	0.03
990211	I	8	0	5	4	-20.0	400	0.03
		11	0	2	7	-20.0	400	0.03
990407	I	22	175	9	1	-25.0	400	0.03
		22	175	1	9	-24.0	400	0.04
991110	I	33	175	5	4	-22.0	400	0.04
		30	175	2	8	-23.1	400	0.04
991111	I	22	175	5	4	-22.0	400	0.04
		23	175	2	8	-23.1	400	0.04
011114	D	21	125	5	4	-22.0	400	0.04
		12	125	2	8	-23.1	400	0.04
011126	D	18	125	5	4	-22.0	400	0.04
		18	125	2	8	-23.1	400	0.04
011205	D	17	125	5	4	-22.0	400	0.04
		17	125	2	8	-23.1	400	0.04

^aI: states are drawn independently across periods; D: states are drawn without replacement from an urn which initially contains 18 balls, 6 for each state.

^bAs discussed in the text, endowment of Notes includes loans to be repaid at the end of the period.

our experimental markets occurs throughout each period, no information is revealed during that time; information is revealed only after trading ends, when the state of nature is drawn.)

Given our focus on end-of-period prices and holdings, it is appropriate to organize the data using a static model of asset trading, as in Arrow and Hahn (1971) or Radner (1972): investors trade assets before the state of nature is known; assets yield dividends and consumption takes place after the state of nature is revealed. (Because there is only one good, there is no trade in commodities, hence no trade after the state of nature is revealed.)

Table 3 summarizes end-of-period prices in all of our experiments. Note that prices are below expected returns in the vast majority of cases; this provides evidence that subjects are indeed strictly risk averse. As mentioned before, this confirms observations in experiments on individual choice; see (Holt and Laury (2002)).

We are interested in the deviation of prices and portfolio choices from CAPM predictions. We use Sharpe ratios to provide a measure of the deviation of the market portfolio from mean-variance efficiency, and mean absolute deviations to provide a measure of the deviation of individual choices from perfect portfolio separation.¹⁶

Recall that, given asset prices q , the *rate of return* on a portfolio θ is $E[\text{div } \theta/q \cdot \theta]$, and the *excess rate of return* is the difference between the return on θ and the 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 $E[\text{div } \theta/q \cdot \theta] - 1$. The *Sharpe ratio* of θ is the ratio of its excess return to its volatility:

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

The market portfolio is mean variance efficient if and only if the market portfolio has the largest Sharpe Ratio among all portfolios, so the difference between the maximum Sharpe ratio of any portfolio and the Sharpe ratio of

¹⁶Other measures of deviation from mean variance efficiency would yield very similar results; in using Sharpe ratios we follow Gibbons, Ross and Shanken (1989).

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

^aSecurity.

^bEnd-of-period transaction price/expected payoff.

^cNotes.

^dFor Notes, end-of-period transaction prices only are displayed. Payoff equals 100.

^eEnd-of-period transaction prices only are displayed. Expected payoffs are as in 981007. Same for 990211, 990407, 991110 and 991111.

the market portfolio

$$\max_{\theta} \text{ShR}(\theta) - \text{ShR}(m)$$

is a measure of the deviation of the market portfolio from mean-variance efficiency.

Because a typical experiment involves more than 30 subjects, displaying portfolio holdings of each subject in each experiment is impractical (and would not be very informative). Instead, we focus on the average deviation between actual holdings of risky assets and the holdings of risky assets predicted by CAPM. Portfolio Separation predicts that each investor's holding of risky assets should be a non-negative multiple of the market portfolio of risky assets; equivalently, that the ratio of the value of investor i 's holding of asset A to the value of investor i 's holding of all risky assets should be the same as the ratio of the value of the market holding of asset A to the value of the market portfolio of all risky assets. A measure of the extent to which the data deviates from the prediction is the mean absolute difference of these ratios:

$$\frac{1}{I} \sum \left| \frac{p_A \theta_A^i}{p \cdot \theta^i} - \frac{p_A m_A}{p \cdot m} \right|$$

Figures 1, 2 and 3 summarize the results from a typical experiment.¹⁷ Figure 1 shows the complete history of prices; note that all transaction prices are below expected payoffs, reflecting the presence of substantial risk aversion. (Expected payoffs for asset A are the higher horizontal lines in each period, and expected payoffs for asset B are the lower horizontal lines; because states were drawn without replacement, expected payoffs are not constant across periods.) Figure 2 shows, with respect to prices at each transaction, the deviation of the market portfolio from mean-variance efficiency; note that within each period and over the course of the experiment, pricing comes to more closely approximate mean variance pricing. Figure 3 shows individual portfolio holdings at the end of each period; note that holdings appear quite random (subject to the accounting identity that holdings sum to the market portfolio).

¹⁷Again: a complete record of every transaction in every experiment is available at <http://www.hss.caltech.edu/~pbs/BPZdata>.

Casual observation of Figures 2 and 3 suggests that approximate mean–variance efficiency appears to prevail (at the end of the period) in at least half the periods, but Portfolio Separation fails spectacularly in *every* period. Figure 4, which summarizes the end-of-period prices and choices in all experiments, makes this point much more sharply. Each point (small circle) in Figure 4 represents a single period of a single experiment. The horizontal component of each point is the deviation of the market portfolio from mean-variance efficiency (at end-of-period prices); the vertical component of each point is the mean absolute deviation from portfolio separation (at end-of-period holdings). As can be seen very clearly in Figure 4, there is *no correlation* between the deviation from mean variance efficiency and the deviation from portfolio separation.¹⁸

5 Econometric Results

5.1 Testing Strategy

The (asymptotic) distribution of the GMM statistic under the CAPM+ ϵ is non-central χ^2 with one degree of freedom (the number of risky assets minus one), and with unknown non-centrality parameter. Our test builds on this property. Specifically, we compute the GMM statistic for the 60+ periods (samples) across our experiments. These outcomes are then used to construct empirical distribution functions of the GMM statistic.

¹⁸The deviations from portfolio separation are particularly striking: the average deviations are roughly as large as they would be if investors chose portfolio weightings *at random*. To make the point most simply, suppose all investors hold the same risky wealth but choose the weighting on asset A at random. The population mean of weighting on asset A must then equal the market weighting on asset A , which is approximately .4 in many of our experiments. This will be the case if weightings on asset A are drawn independently from the distribution

$$\frac{3}{2}\lambda_{[0,.4]} + \frac{2}{3}\lambda_{[.6,1]}$$

where λ_E denotes the restriction of Lebesgue measure to $E \subset [0, 1]$. In that case, the mean absolute deviation will be only .24.




Figure 1: History of transaction prices in experiment 011126




Figure 2: Deviation from Mean Variance Efficiency computed with respect to transaction prices in experiment 011126




Figure 3: Individual end-of period portfolio holdings in experiment 011126




Figure 4: 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 transaction prices), all periods in all experiments. There is no correlation between distance from CAPM pricing (x -axis) and violation of portfolio separation (y -axis).

We cannot readily aggregate the results over all periods, or even across all periods in a single experiment, because periods within a single experiment are not independent (subject populations are the same), so that the GMM statistics across periods within an experiment are not independent. However, periods in *different* experiments are independent (because the subject populations of different experiments are disjoint). We therefore construct 8 samples; the first sample consists of first periods of all experiments, the second sample consists of second periods of all experiments, and so forth. We then test whether the empirical distribution of GMM statistics is non-central χ^2 for each sample. (We use only the first 8 periods because we have only a few experiments which last 9 periods.)

We use both the Kolmogorov-Smirnov statistic and the Cramer-Von Mises statistic. The former uses the supremum of the deviations of the empirical distribution function (of the GMM statistic) from a non-central χ^2 distribution function; the latter uses the density-weighted mean squares of these deviations. We estimate the non-centrality parameter from all the data (all periods in all experiments) in order to minimize estimation error. Effectively, the non-centrality parameter is estimated on the basis of a sample that is *at least seven times as large* as the samples on which we test whether the empirical distribution function of the GMM statistic is non-central χ^2 .¹⁹

5.2 Test Results

Figure 5 depicts the empirical distribution of the logarithm of the GMM statistic across all periods in our experiments. The smooth line is the distribution of the logarithm of a *central* χ^2 -distributed random variable; the jagged line is the empirical distribution of our test statistic. As can be seen, the jagged line appears to be a horizontal translation of the smooth line, which suggests that the GMM statistics are drawn from a non-central χ^2 distribution. This suggestion is confirmed in Table 4, which reports the findings

¹⁹An alternative approach would be to estimate the non-centrality parameter in-sample and adjust p values accordingly. We have not done this because the correct adjustments are not known.

of the Kolmogorov-Smirnov (KS) and Cramer-von Mises (CvM) tests applied to our model. For each sample, we test the fit of the empirical distribution function of our GMM statistics to a non-central χ^2 distribution, using the best fit for the unknown non-centrality parameter (11.6 for for KS, 10.0 for CvM).²⁰ In Table 4 we follow Shorack and Wellner (1986, p. 239) to correct for small-sample biases; p values are obtained from the same source. (Critical values for the Cramer-von Mises statistic are known only for some p values; for other p -values, we report a range.)

At the 1% level, both KS and CvM goodness-of-fit tests reject only in period 2, but fail to reject in other periods. At the 5% level, both KS and CvM reject only in periods 1 and 2, but fail to reject in other periods. It might be useful to keep in mind that econometric tests of asset pricing models against historical data frequently reject at much smaller values of p , and that our tests are more demanding because they test prices *and* holdings. For example, in arguing that the performance of the three-factor model is superior to other models, despite the fact that it is rejected at the 0.5% significance level, Davis, Fama and French (2000, p. 450) write : “[...] the three-factor model [...] is rejected by the [...] test. This result shows that the three-factor model is just a model and thus an incomplete description of expected returns. What the remaining tests say is that the model’s shortcomings are just not those predicted by the characteristics model.”

We have argued that the correlation between prices and perturbations inherent in CAPM+ ϵ means that the correct comparison of the distribution of our test statistic is with a *non-central* χ^2 distribution. Tables 5 and 6 show the result of ignoring this correlation and comparing the distribution of our test statistic with a *central* χ^2 distribution. (We separate experiments in which states were drawn with replacement from experiments where states

²⁰Best fits are obtained as follows. Let $F_E(\cdot)$ denote the empirical distribution function of the GMM statistic. Let $F_{\lambda^2}(\cdot)$ denote the χ^2 distribution with one degree of freedom and non-centrality parameter λ^2 . The best fit is obtained as

$$\inf_{\lambda^2} \sup_x |F_E(x) - F_{\lambda^2}(x)|.$$




Figure 5: Empirical distribution of the GMM statistic, all periods in all experiments (jagged line), against a *central* χ^2 distribution. A noncentral χ^2 distribution provides a better fit, consistent with the small-sample biases expected if CAPM+ ϵ is correct.

Table 4: Tests Of CAPM+ ϵ Accommodating Correlation Between Prices And Perturbations

Period Number	Number of Observations	KS ^a	p value ^b	CvM ^c	p value ^d
1	9	1.53	$0.05 > p > 0.025$	0.49	$0.05 > p > 0.025$
2	9	2.01	$p < 0.01$	0.91	$p < 0.01$
3	9	1.01	$p > 0.15$	0.21	$p > 0.15$
4	9	1.33	$0.15 > p > 0.10$	0.30	$0.15 > p > 0.10$
5	8	1.26	$0.10 > p > 0.05$	0.21	$p > 0.15$
6	9	0.80	$p > 0.15$	0.21	$0.10 > p > 0.05$
7	6	0.87	$p > 0.15$	0.04	$p > 0.15$
8	4	1.26	$0.10 > p > 0.05$	0.42	$0.10 > p > 0.05$

^aKolmogorov-Smirnov (KS) statistic of the difference between the empirical distribution function of GMM statistics across experiments for a fixed period and a non-central χ^2 distribution with non-centrality parameter 11.6. The KS statistic is modified for small sample bias. See Shorack and Wellner (1986) [p. 239].

^bBased on Table 1 on p. 239 of Shorack and Wellner (1986).

^cCramer-von Mises (CvM) statistic of the difference between the empirical distribution function of GMM statistics across experiments for a fixed period and a non-central χ^2 distribution with non-centrality parameter 10.0. The CvM statistic is modified for small sample bias. See Shorack and Wellner (1986) [p. 239].

^dBased on Table 1 on p. 239 of Shorack and Wellner (1986).

were drawn without replacement only because a combined table would be too large to display legibly on a single page.)

5.3 Power

There are several reasons to believe our tests have substantial power.

- i) We require the non-centrality parameter to be the same across periods *as well as across experimental sessions*. Since distributional properties of the individual perturbation terms ultimately determine the value of the non-centrality parameter, this means that we implicitly assume that these properties do not change across experiments. In other words, we impose a strong homogeneity assumption across different subject populations.
- ii) The non-centrality parameter imposes a *tight relationship between the moments of the GMM statistic*. In particular, the difference between its variance and its mean is equal to the (fixed) number of degrees of freedom plus three times the non-centrality parameter.

Formal analyses of power (or size) are often accomplished by means of Monte Carlo analysis. In such analysis, values are posited for all the parameters of one's economic model — in particular, for the parameters that determine the distributions of the random variables — and draws from these distributions are used to generate a sequence of samples of a size comparable to the size of the actual sample(s) in the empirical study. The sequence of samples induces a sequence of statistics, with which one can construct a distribution — an estimate of the finite-sample distribution under the maintained parametric assumptions. Typically this exercise is conducted under a parametrization for which the null hypothesis is satisfied, which provides an estimate of the true size of one's test (true probability of rejecting a correct null based on the actual cut-offs used in the empirical test), and then again under various parametrizations for which the null hypothesis is violated, in order to trace a power function (the probability of rejecting the null when it is false, as a function of changes in such parametrizations).

Table 5: GMM Tests Of CAPM+ ϵ Ignoring Correlation Between Prices and Perturbations — Experiments where Draws were Independent

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 for $\lambda = 0$.00	.14	.00	.00	.00	.00			
	β^I ($*10^{-3}$)	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 for $\lambda = 0$.00	.35	.32	.04	.06	.00			
	β^I ($*10^{-3}$)	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 for $\lambda = 0$.02	.00	.02	.00	.04	.00	.69		
	β^I ($*10^{-3}$)	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	0.6	12.5	103.1	† ^b	0.2	13.6	†	
	p level for $\lambda = 0$.02	.44	.00	.00	-	.66	.00	-	
	β^I ($*10^{-3}$)	0.5	0.5	0.3	-0.3	-	0.3	0.2	-	
	s.e. ($*10^{-3}$)	0.0	0.0	0.0	0.1	-	0.1	0.1	-	
991110	χ_1^2	197.5	72.8	31.6	7.4	2.7	6.9			
	p level for $\lambda = 0$.00	.00	.00	.01	.10	.01			
	β^I ($*10^{-3}$)	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 for $\lambda = 0$.03	.23	.00	.00	.00	.00	.00	.00	
	β^I ($*10^{-3}$)	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	

^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 6: GMM Tests Of CAPM+ ϵ Ignoring Correlation Between Prices and Perturbations — Experiments where Draws were without Replacement

Experiment	Statistic	Periods								
		1	2	3	4	5	6	7	8	9
011114	χ_1^2	37.8	2.3	2.2	5.5	18.2	15.4	10.7	22.1	4.3
	p level for $\lambda = 0$.00	.13	.13	.02	.00	.00	.00	.00	.04
	β^I ($*10^{-3}$)	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 for $\lambda = 0$.00	.44	.01	.02	.23	.25	.65	.01	
	β^I ($*10^{-3}$)	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	10.8	18.2	13.9	13.6	15.6	7.9	19.0	5.0
	p level for $\lambda = 0$.09	.00	.00	.00	.00	.00	.01	.00	.02
	β^I ($*10^{-3}$)	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	0.1	0.1

In our context, however, Monte Carlo power analysis seems both infeasible and arbitrary.

- Monte Carlo analysis seems infeasible in our setting because the relevant parameter space is infinite-dimensional. In our context, the relevant parameters are the distributions of endowments, mean-variance demand functions and perturbations. Endowments and mean-variance demand functions are determined by four real parameters, but perturbations are *functions of unknown form*. Hence, the perturbations are specified by a distribution on an infinite-dimensional function space. One might restrict attention to a particular class of perturbation functions — but we have no idea what class could/should be used.
- Monte Carlo seems arbitrary in our setting because power analysis requires specification of plausible parametrizations for which the null hypothesis is violated. In our setting, there are *many* ways in which the null hypothesis could be violated, and it is not at all obvious what possible violations should be thought most relevant; obvious choices

include the possibility that perturbation functions are not drawn independently, or that they are not drawn from a distribution with mean zero violations, or that the market is not in equilibrium at the end of each experimental period.

One might respond to these objections by using the data to suggest parametrizations and relevant violations of the null hypothesis. However, a much more straightforward approach to power analysis is available. Monte Carlo analysis is most useful when one does not have data — but we do have data, generated by actual economies. It seems more natural therefore to *use the data we have* to generate tests of power.²¹

Motivated by this reasoning, we offer two tests of power generated with the existing data:

- In the theory, choices determine prices. Our test does not reject the theory; i.e., data appear to be consistent with the theory. Asking for the power of the test amounts to asking whether the tests would have rejected the theory on *other* data. To answer this question, we *systematically perturb prices in 3 randomly chosen experiments*. The new prices are within 5 francs (2.5%) of the prices found in the experiments; higher prices for security A and lower prices for security B. In comparison to the range of prices we observe in our experiments — both across experiments and within a single experiment — this is a marginal change in prices. To put it in statistical terms: the unconditional probability (i.e., the probability not conditional on subjects' actual choices) of observing such prices is high. As can be seen in Table 6, if we apply our tests to the artificial data (perturbed prices, same choices), we reject the theory at the 5% level in 50% of all samples for both the KS and CvM tests, whereas the on the actual data we reject at the 5% level in only 25% of all samples. Put differently, the theory would be strongly rejected if as few as 1/3 of the prices differed from

²¹Of course, the use of one's own data also underlies bootstrap estimation, which has replaced Monte Carlo analysis to estimate the correct size of one's test under the null hypothesis.

actual prices by as little as 2.5%; this provides an indirect demonstration of substantial power. Thus, one cannot explain actual choices by prices that are only slightly different — even if those prices are likely to be observed in *other periods*. Thus, there is something special about the prices in a given period: they match the choices in that period *and only in that period*.²²

- Our second exercise investigates power against violations of the statistical assumption that economies are drawn independently. In order to assure independence, we used samples consisting of the outcomes (prices; choices) in a (fixed) period across all experiments; this guarantees independence because subject populations are disjoint across experiments. If, instead, we use outcomes from all periods of a *given* experiment, we would then have samples that fail the assumption of independence, because subject populations are constant within a single experiment, whence choices within an experiment — which determine prices — are presumably correlated as well. When we apply our goodness-of-fit tests to these experiment-wise samples, our p -values go down significantly and we reject more often. As Table 7 shows, at the 5% level the KS test now rejects 44% of the time and the CvM test rejects 33% of the time. The decrease in p -values and increase in number of rejections demonstrates power for purely statistical violations of our null hypothesis.

6 Conclusion

In this paper, we provide a rationale for testing asset pricing models that rely on portfolio separation — such as CAPM and its multi-factor extensions —

²²An alternative would have been to systematically perturb prices in *all* experiments. However, in that case we would have changed the entire (unconditional) distribution of prices, and the distribution would not have been the one obtained in the experiments. Hence such a price perturbation seems to be beside the point. We think that our exercise is the right one.

Table 7: Power Analysis: Tests Of CAPM+ ϵ Accommodating Correlation Between Prices And Perturbations – Tests Per Experiment (Violating Independence)

Period Number	Number of Observations	KS ^a	p value ^b	CvM ^c	p value ^d
981007	6	1.43	$0.05 > p > 0.025$	0.70	$0.025 > p > 0.01$
981116	6	1.50	$0.025 > p > 0.01$	0.41	$0.10 > p > 0.05$
990211	7	1.20	$0.15 > p > 0.10$	0.10	$p > 0.15$
990407	6	0.86	$p > 0.15$	0.04	$p > 0.15$
991110	6	1.27	$0.10 > p > 0.05$	0.23	$p > 0.15$
991111	8	2.19	$p < 0.01$	1.19	$p < 0.01$
011114	9	0.95	$p > 0.15$	0.10	$p > 0.15$
011126	8	1.81	$p < 0.01$	0.83	$p < 0.01$
011205	9	0.54	$p > 0.15$	0.06	$p > 0.15$

^aKolmogorov-Smirnov (KS) statistic of the difference between the empirical distribution function of GMM statistics across experiments for a fixed period and a non-central χ^2 distribution with non-centrality parameter 11.6. The KS statistic is modified for small sample bias. See Shorack and Wellner (1986) [p. 239].

^bBased on Table 1 on p. 239 of Shorack and Wellner (1986).

^cCramer-von Mises (CvM) statistic of the difference between the empirical distribution function of GMM statistics across experiments for a fixed period and a non-central χ^2 distribution with non-centrality parameter 10.0. The CvM statistic is modified for small sample bias. See Shorack and Wellner (1986) [p. 239].

^dBased on Table 1 on p. 239 of Shorack and Wellner (1986).

even in the absence of convincing evidence for such portfolio separation. We offer a theoretical model, a novel econometric procedure to test the model, and tests based on data from experimental financial markets. These tests fail to reject the model.

Our analysis suggests several lessons. The first is that, despite the modest risks, experimental financial markets can provide significant and useful insights. A second is that the standard model of choice under unobserved heterogeneity that is widely used applied economics should be used with some care. In finite markets, unexplained heterogeneity in demands (usually a key determinant of the unexplained portion of observed choices) need *not* be orthogonal to prices, and this may have significant effects on the econometric analysis.

In contrast to standard econometric analysis, the econometric procedure introduced here explicitly links prices and choices. We have applied this procedure only to data from experimental financial markets, but it is applicable to historical data as well — provided suitable choice data can be found.

Appendix A: CAPM

To derive the conclusions of CAPM in our setting in which short sales of risky assets are not permitted, we begin by analyzing the setting in which arbitrary short sales *are* permitted. Write $\widehat{Z}^i(p)$, respectively $\widehat{z}^i(p)$, for investor i 's demand for all assets, respectively risky assets, when the price of risky assets is p . Note that either of $\widehat{Z}^i(p)$, $\widehat{z}^i(p)$ determines the other (through budget balance); we will focus on whichever is convenient for the purpose at hand.

Assuming, as we do throughout, that consumptions are in the range where preferences are monotone, the first-order conditions for optimality and some algebra show that

$$\widehat{z}^i(p) = \frac{1}{b^i} \Delta^{-1}(\mu - p) \quad (15)$$

At equilibrium, the demands for risky assets must clear the market for risky assets, so if \widehat{p} is an equilibrium price then:

$$\sum_{i=1}^I \widehat{z}^i(\widehat{p}) = m$$

From these equations we can solve for the unique equilibrium price \widehat{p} :

$$\widehat{p} = \mu - \left(\sum_{i=1}^I \frac{1}{b^i} \right)^{-1} \Delta m = \mu - \left(\frac{1}{I} \sum_{i=1}^I \frac{1}{b^i} \right)^{-1} \Delta \bar{m}$$

In our setting, short sales of risky assets are not permitted, and demand functions are *not* given by the equation (15). However, we assert that the model with short sales and the our model without short sales admit the same equilibrium prices.

To see this, write $z^i(p)$ for investor i 's demand for risky assets when prices are p and short sales of risky assets are not permitted. Note that $z^i(p) = \widehat{z}^i(p)$ whenever $\widehat{z}^i(p) \geq 0$: in particular, $z^i(\widehat{p}) = \widehat{z}^i(\widehat{p})$; it follows immediately that \widehat{p} is an equilibrium price in the setting when short sales of risky assets are not permitted. To see that there is no *other* equilibrium price in this setting, suppose that $p^* \neq \widehat{p}$ were such an equilibrium price. If constrained demand $z^j(p^*)$ were strictly positive for some investor j , then

constrained demand $z^j(p^*)$ would coincide with unconstrained demand $\hat{z}^j(p^*)$ for investor j . However, formula (15) guarantees that if $\hat{z}^j(p^*)$ were positive for some investor j then $\hat{z}^i(p^*)$ would be positive for every investor i , whence $z^i(p^*)$ would coincide with $\hat{z}^i(p^*)$ for every investor i . Because $p^* \neq \hat{p}$, this would imply that p^* was not an equilibrium price after all. It follows that constrained demand $z^j(p^*)$ cannot be strictly positive for any investor j . At equilibrium, asset markets clear; because the market portfolio is strictly positive, it follows that some investor k chooses an equilibrium portfolio that involves the risky asset A but not the risky asset B and some investor ℓ chooses an equilibrium portfolio that involves the risky asset B but not the risky asset A :

$$\begin{aligned} z_A^k(p^*) &> 0 \quad , \quad z_B^k(p^*) = 0 \\ z_A^\ell(p^*) &= 0 \quad , \quad z_B^\ell(p^*) > 0 \end{aligned}$$

The first-order conditions for investors k, ℓ entail:

$$\begin{aligned} \frac{p_A^*}{p_B^*} &\leq \frac{\text{MU}_A^k}{\text{MU}_B^k} \\ \frac{p_B^*}{p_A^*} &\leq \frac{\text{MU}_B^\ell}{\text{MU}_A^\ell} \end{aligned}$$

Direct calculation using the explicit form of utility functions and making use of the fact that $\text{var}(x + y) = \text{var}(x) + 2\text{cov}(x, y) + \text{var}(y)$ yields

$$\begin{aligned} \frac{\text{MU}_A^k}{\text{MU}_B^k} &= \frac{E(A) - b^k z_A^k(p^*) \text{var}(A)}{E(B) - b^k z_A^k(p^*) \text{cov}(A, B)} \\ \frac{\text{MU}_B^\ell}{\text{MU}_A^\ell} &= \frac{E(B) - b^\ell z_B^\ell(p^*) \text{var}(B)}{E(A) - b^\ell z_B^\ell(p^*) \text{cov}(A, B)} \end{aligned}$$

Our assumptions guarantee that $b^k, b^\ell, \text{var}(A), \text{var}(B)$ are all strictly positive, and the particular structure of payoffs of the risky assets and state probabilities guarantee that $\text{cov}(A, B) < 0$. Combining all these yields:

$$\begin{aligned} \frac{p_A^*}{p_B^*} &\leq \frac{\text{MU}_A^k}{\text{MU}_B^k} < 1 \\ \frac{p_B^*}{p_A^*} &\leq \frac{\text{MU}_B^\ell}{\text{MU}_A^\ell} < 1 \end{aligned}$$

This is a contradiction, so we conclude that \hat{p} is the unique equilibrium price, as asserted.

Appendix B: CAPM $_{+\epsilon}$

In this Appendix we give a formal and rigorous presentation of the idea of the true economy as drawn from a distribution of individual characteristics and as a perturbation of a mean-variance economy. Although the ideas are simple, and the conclusions both intuitive and expected, the details require a little care. As we shall see, the analysis depends on the fact that $\text{cov}(A, B) < 0$ (which guarantees that the equilibria of the ideal economy with short sales and without short sales are the same) and on the fact that the CAPM equilibrium of the ideal economy is regular (which guarantees that the perturbed economy has an equilibrium near the CAPM equilibrium of the ideal economy).

We work throughout in the setting of Section 2 and Appendix A, and retain the same notation. In particular, two risky assets and one riskless asset are traded; the risky assets cannot be sold short but the riskless asset can be; the covariance of the risky assets is negative; consumption may be negative. We normalize throughout so that the price of the riskless asset is 1; the vector of asset prices is $q = (p, 1) \in R_{++}^3$. As before, we write μ for the vector of expected returns on risky assets and Δ for the covariance matrix of risky assets.

Distributions and Draws from a Distribution

We follow Hart, Hildenbrand and Kohlberg (1979) in describing economies as distributions on the space of investor characteristics. The usual description of an investor is in terms of an endowment bundle of commodities and preferences over commodity bundles, but we find it more convenient to adopt a description in terms of an endowment portfolio of assets and a demand function for assets. We assume that endowments and prices, hence wealth, lie in

given compact sets $End \subset \mathbf{R}_+^2 \times \mathbf{R}$, $\mathcal{P} \subset \mathbf{R}_{++}^3$, $[0, \bar{w}] \subset \mathbf{R}_+$. An investor is characterized by an endowment $\omega \in End$ of riskless and risky assets and by a continuous demand function

$$Z : \mathcal{P} \times [0, \bar{w}] \rightarrow \mathbf{R}_+^2 \times \mathbf{R}$$

for risky and riskless assets as a function of wealth $w \in [0, \bar{w}] \subset \mathbf{R}_+$ and prices for risky assets $p \in \mathcal{P} \subset \mathbf{R}_{++}^2$. (Recall that we have normalized so that the price of the riskless asset is 1.) We assume throughout that the value of demand is equal to the value of the endowment:

$$(p, 1) \cdot Z(p, (p, 1) \cdot \omega) = (p, 1) \cdot \omega \leq [0, \bar{w}]$$

for each ω, p . (We could assume that demand satisfies properties that follow from revealed preference, but there is no need to do so.) Write \mathcal{D} for the space of demand functions, and equip \mathcal{D} with the topology of uniform convergence. $End \times \mathcal{D}$ is the space of *investor characteristics*.

We view a compactly supported probability measure τ on $End \times \mathcal{D}$ as the distribution of investor characteristics in a fixed economy and also as the distribution of characteristics of the pool from which economies are drawn. Given an integer I , a particular draw of I investors from τ can be described by a distribution of the form

$$\tilde{\tau} = \frac{1}{I} \sum_{i=1}^I \delta_{(\omega^i, Z^i)}$$

where $\delta_{(\omega^i, Z^i)}$ is point mass at the characteristic $(\omega^i, Z^i) \in \text{supp } \tau \subset End \times \mathcal{D}$. We identify the set of such draws with $(\text{supp } \tau)^I$, which, by abuse of notation, we view as a subset of $\mathcal{M}(E \times \mathcal{D})$, the space of all compactly supported probability measures on $End \times \mathcal{D}$. The I -fold product measure τ^I on $(\text{supp } \tau)^I$ is the distribution of all draws.

Equilibrium

Given a distribution $\eta \in \mathcal{M}(E \times \mathcal{D})$, an *equilibrium* for η is a price $p \in \mathcal{P}$ such that

$$\int Z(p, (p, 1) \cdot \omega) d\eta = \int \omega d\eta$$

(Because we describe investor characteristics in terms of demand functions, we focus on prices and suppress consumptions. Of course, $Z(p, (p, 1) \cdot \omega)$ is the equilibrium consumption of the investor with characteristics (ω, Z) .)

We caution the reader that a distribution η need not admit an equilibrium, and that convergence of distributions does *not* imply convergence of (sets of) equilibria. However, as we shall show, the situation is much better for the distributions of most interest to us.

CAPM Distributions

Given an endowment ω , the portfolio θ is *budget feasible* if $(p, 1) \cdot \theta \leq (p, 1) \cdot \omega$ for every $p \in \mathcal{P}$. We say $\sigma \in \mathcal{M}(\text{End} \times \mathcal{D})$ is a *mean-variance distribution* if for each $(\omega, Z) \in \text{supp } \sigma$ there is a coefficient of risk aversion $b(\omega, Z) > 0$ such that the mean-variance utility function

$$U^{b(\omega, Z)} = E(x) - \frac{1}{b(\omega, Z)} \text{var}(x)$$

is strictly monotone on the set of dividends of feasible portfolios and Z is the (restriction of) the portfolio demand function derived from $U^{b(\omega, Z)}$. (Keep in mind that we require holdings of risky assets to be non-negative.) Given a mean-variance distribution σ we write $B_\sigma = \int b(\omega, Z)^{-1} d\sigma$ for the *market risk tolerance* and $\bar{m}_\sigma = \int \omega d\sigma$ for the *per capita market portfolio*. We say the mean-variance distribution σ is a *CAPM distribution* if the price

$$p_\sigma = \mu - B_\sigma^{-1} \Delta \bar{m}_\sigma$$

belongs to the interior $\text{int } \mathcal{P}$ of \mathcal{P} . If σ is a CAPM distribution, it follows as in Appendix A that σ admits p_σ as the *unique* equilibrium price, that the mean market portfolio \bar{m} is mean-variance efficient at prices p_σ and that equilibrium holdings of risky assets $z(p_\sigma, (p_\sigma, 1) \cdot \omega)$ are non-negative multiples of the mean market portfolio \bar{m} (portfolio separation).

Mean Zero Perturbations

Write $\pi_E : E \times \mathcal{D} \rightarrow E$ for the projection on the first factor. If $\tau, \sigma \in \mathcal{M}(E \times \mathcal{D})$ we say τ is a *perturbation* of σ if there is a measurable function $f : \text{supp } \tau \rightarrow \mathcal{D}$ such that $\sigma = (\pi_E, f)_* \tau$; that is,

$$\sigma(B) = \tau((\pi_E, f)^{-1}(B)) = \tau\{(\omega, Z) : (\omega, f(\omega, Z)) \in B\}$$

for each Borel set $B \subset E \times \mathcal{D}$. We say τ is a *mean zero* perturbation of σ if in addition

$$\int [Z - f(\omega, Z)] d\tau = \int Z d\tau - \int Z' d\sigma = 0$$

Evidently, if τ is a mean zero perturbation of σ , then τ and σ admit the same equilibria — although neither may admit any equilibrium at all.

Perturbations of CAPM Distributions

We are now in a position to state and prove the result we require.

Theorem *Let σ be a CAPM distribution, and let τ be a mean-zero perturbation of σ . For each $\varepsilon_0 > 0$ there is an integer I_0 and for every $I > I_0$ there is a subset $\Gamma_I \subset (\text{supp } \tau)^I$ such that*

- i) $\tau^I(\Gamma_I) > 1 - \varepsilon_0$*
- ii) for every $\gamma \in \Gamma_I$, the draw $\tilde{\gamma} = F(\gamma)$ from σ admits a unique equilibrium $p_{\tilde{\gamma}}$ and the draw γ from τ admits at least one equilibrium*
- iii) if $\gamma \in \Gamma_I$ and p_γ is any equilibrium of γ then $\|p_\gamma - p_{\tilde{\gamma}}\| < \varepsilon_0$*

Informally: if we draw a large enough sample from τ then, with high probability the sample economy and the CAPM economy of which it is a perturbation have nearly the same equilibrium price(s).

Proof If ν is a distribution, let $D_\nu : \mathcal{P} \rightarrow \mathbf{R}_+^2 \times \mathbf{R}$ be the market demand function for assets:

$$D_\nu(p) = \int Z(p, (p, 1) \cdot \omega) d\nu$$

and let \bar{M}_ν be the per capita market portfolio

$$\bar{M}_\nu = \int \omega d\nu$$

By definition, an equilibrium for ν is a zero of excess demand $D_\nu - \bar{M}_\nu$.

By assumption, $p_\sigma \in \text{int } \mathcal{P}$. Choose $\varepsilon_1 < \varepsilon_0$ so that $B(p_\sigma, \varepsilon_1) \subset \mathcal{P}$. Direct computation shows that the excess demand function $D_\sigma - \bar{M}_\sigma$ is regular at p_σ . It follows that there is an $\varepsilon_2 > 0$ such that if $H : \mathcal{P} \rightarrow \mathbf{R}_+^2 \times \mathbf{R}$ is any continuous function and

$$\|H - (D_\sigma - \bar{M}_\sigma)\|_{B(p_\sigma, \varepsilon_1)} < \varepsilon_2$$

then H has at least one zero in $B(p_\sigma, \varepsilon_1)$.

On the other hand, $D_\sigma - \bar{M}_\sigma$ is bounded away from 0 on $\mathcal{P} \setminus B(p_\sigma, \varepsilon_1)$, so there is an $\varepsilon_3 > 0$ such that if

$$\|H - (D_\sigma - \bar{M}_\sigma)\|_{\mathcal{P} \setminus B(p_\sigma, \varepsilon_1)} < \varepsilon_3$$

then H is bounded away from 0 on $\mathcal{P} \setminus B(p_\sigma, \varepsilon_1)$. Setting $\varepsilon = \min\{\varepsilon_1, \varepsilon_2, \varepsilon_3\}$, we conclude: if $H : \mathcal{P} \rightarrow \mathbf{R}_+^2 \times \mathbf{R}$ is any continuous function for which

$$\|H - (D_\sigma - \bar{M}_\sigma)\|_{\mathcal{P}} < \varepsilon$$

then H has at least one zero on \mathcal{P} , and all its zeroes belong to $B(p_\sigma, \varepsilon)$, and hence to $B(p_\sigma, \varepsilon_0)$.

For each I , set

$$\begin{aligned} \mathcal{G}_I^1 &= \{\gamma \in (\text{supp } \tau)^I : \|Z_\gamma - Z_\tau\| < \varepsilon/2\} \\ \mathcal{H}_I^1 &= \{\gamma \in (\text{supp } \tau)^I : \|\bar{M}_\gamma - \bar{M}_\tau\| < \varepsilon/2\} \\ \mathcal{G}_I^2 &= \{\zeta \in (\text{supp } \sigma)^I : \|Z_\zeta - Z_\tau\| < \varepsilon/2\} \\ \mathcal{H}_I^2 &= \{\zeta \in (\text{supp } \sigma)^I : \|\bar{M}_\zeta - \bar{M}_\sigma\| < \varepsilon/2\} \\ \Gamma_I &= \mathcal{G}_I^1 \cap \mathcal{H}_I^2 \cap F^{-1}(\mathcal{G}_I^2) \cap F^{-1}(\mathcal{H}_I^2) \end{aligned}$$

Market demand is the expectation of individual demand, and the per capita market portfolio is the expectation of individual endowment portfolios, so applying the Strong Law of Large Numbers in the space of continuous functions $\Phi : \mathcal{P} \rightarrow \mathbf{R}^3$ (see Ledoux and Talagrand (1991) for the appropriate Banach space version) and in \mathbf{R}^3 implies that there is an index I_0 such that if $I > I_0$ then

$$\begin{aligned}\tau^I(\mathcal{G}_I^i) &> 1 - \frac{\varepsilon}{4} \\ \sigma^I(\mathcal{H}_I^i) &> 1 - \frac{\varepsilon}{4}\end{aligned}$$

for $i = 1, 2$.

Let f be the function given in the definition of mean zero perturbation, and write $F = (\pi_E, f)$. By assumption

$$\tau(F^{-1}(\mathcal{G}_I^2)) = \sigma(\mathcal{G}_I^2) \quad \text{and} \quad \tau(F^{-1}(\mathcal{H}_I^2)) = \sigma(\mathcal{H}_I^2)$$

so if $I > I_0$ then $\tau(\Gamma_I) > 1 - \varepsilon$.

Finally, if $\gamma \in \mathcal{G}_I$ then

$$\begin{aligned}\|(D_\gamma - \overline{M}_\gamma) - (Z_\gamma - \overline{M}_\gamma)\| &< \varepsilon \\ \|(D_{F_*\gamma} - \overline{M}_{F_*\gamma}) - (Z_\sigma - \overline{M}_\sigma)\| &< \varepsilon\end{aligned}$$

Our construction guarantees that γ and $F_*\gamma$ each admit at least one equilibrium and that all these equilibria lie in $B(p_\sigma, \varepsilon_0)$. Finally, because $F_*\gamma$ is a CAPM economy, it actually admits a unique equilibrium, so the proof is complete. ■

Appendix C: Estimation of W

We first specify our estimator Ξ^I of W . After that, we provide an asymptotically unbiased and uncorrelated estimator of individual risk tolerances, to be used in the formulation of W . Third, we prove that the error of this estimator does not affect the asymptotic properties of Ξ^I . As a result, we

substitute true risk tolerances for estimates of the risk tolerances in the formula of Ξ^I , and we proceed to the fourth step, where we prove convergence of Ξ^I to W .

In the sequel, we take the risk aversion coefficients as fixed. This is consistent with the theory as long as perturbations are drawn independently from risk aversion coefficients. The econometrics conditions on risk aversion.

Likewise, we assume that individual perturbations are independent across periods within an experimental session. In fact, all we need is that they are asymptotically orthogonal conditional on prices:

$$E\{[\epsilon_t^i]_k [\epsilon_\tau^i]_j | p_1^N, \dots, p_t^I\} \rightarrow 0, \quad (16)$$

all $\tau \neq t$, as $N \rightarrow \infty$.

The Estimator Ξ^I

To understand our estimator Ξ^I of W , let β_t^i denote agent i 's risk tolerance, i.e., $\beta_t^i = 1/b_t^i$. Define the cross-sectional average holding: let

$$\bar{m}^I = \frac{1}{I} \sum_{i=1}^I z_t^i. \quad (17)$$

Also, define:

$$W^I = IE[h_t^I(B_t^I)h_t^I(B_t^I)^\top | p_t^I].$$

So, W is the limit of W^I as $I \rightarrow \infty$. Now re-formulate W^I :

$$\begin{aligned} W^I &= IE[h_t^I(B_t^I)h_t^I(B_t^I)^\top | p_t^I] \\ &= (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i \epsilon_t^i{}^\top | p_t^I] \\ &= (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\epsilon_t^i + \beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) (\epsilon_t^i + \beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\ &\quad - \frac{1}{2} (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\epsilon_t^i + \beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) (E[z_t^i | p_t^I] - \bar{m}^I)^\top | p_t^I] \end{aligned}$$

$$\begin{aligned}
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) (\epsilon_t^i + \beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i (\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) \epsilon_t^{i\top} | p_t^I] \\
= & (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(z_t^i - \bar{m}^I) (z_t^i - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(z_t^i - \bar{m}^I) (\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) (z_t^i - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i (\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\
& -\frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I) \epsilon_t^{i\top} | p_t^I],
\end{aligned}$$

where the last equality follows from

$$\epsilon_t^i = z_t^i - \beta_t^i \Delta^{-1}(\mu - p_t^I).$$

Now consider:

$$\begin{aligned}
\bar{m}^I - \frac{1}{I} \sum_{i=1}^I \beta_t^i \Delta^{-1}(\mu - p_t^I) &= \frac{1}{I} \sum_{i=1}^I (z_t^i - \beta_t^i \Delta^{-1}(\mu - p_t^I)) \\
&= \frac{1}{I} \sum_{i=1}^I \epsilon_t^i,
\end{aligned}$$

which converges to zero, by the law of large numbers. As a result, the second-to-last term of the above expression becomes:

$$\begin{aligned}
& \frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i (\beta_t^i \Delta^{-1}(\mu - p_t^I) - \bar{m}^I)^\top | p_t^I] \\
= & \frac{1}{2}(B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[\epsilon_t^i \left(\beta_t^i \Delta^{-1}(\mu - p_t^I) - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu \Delta^{-1}(\mu - p_t^I) \right)^\top | p_t^I],
\end{aligned}$$

which converges to zero, because $E[\epsilon_t^i | p_t^I] \rightarrow 0$.²³ The same applies to the last term in the above expression.

We shall make the same substitution for \bar{m}^I in the second and third term.

Consequently, there is no difference asymptotically if we define W^I as follows:

$$\begin{aligned} W^I &= (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(z_t^i - \bar{m}^I) (z_t^i - \bar{m}^I)^\top | p_t^I] \\ &\quad - \frac{1}{2} (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E[(z_t^i - \bar{m}^I) \left([\beta_t^i - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu] \Delta^{-1} (\mu - p_t^I) \right)^\top | p_t^I] \\ &\quad - \frac{1}{2} (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I E\left[\left([\beta_t^i - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu] \Delta^{-1} (\mu - p_t^I) \right) (z_t^i - \bar{m}^I)^\top | p_t^I \right]. \end{aligned}$$

Note that this expression does not involve unobservables – except for the risk tolerances β_t^i which we will discuss below.

This suggests the following estimator. Define the cross-sectional covariance of choices, $\text{cov}(z_t^i)$:

$$\text{cov}(z_t^i) = \frac{1}{I} \sum_{i=1}^I (z_t^i - \bar{m}^I) (z_t^i - \bar{m}^I)^\top.$$

Then let

$$\begin{aligned} \Xi^I &= (B_t^I)^2 \text{cov}(z_t^i) \\ &\quad - \frac{1}{2} (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I (z_t^i - \bar{m}^I) \left([\beta_t^i - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu] \Delta^{-1} (\mu - p_t^I) \right)^\top \\ &\quad - \frac{1}{2} (B_t^I)^2 \frac{1}{I} \sum_{i=1}^I \left([\beta_t^i - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu] \Delta^{-1} (\mu - p_t^I) \right) (z_t^i - \bar{m}^I)^\top. \quad (18) \end{aligned}$$

Estimating Risk Tolerances

In order to implement Ξ^I , we need an estimator for the risk tolerances β_t^i . A judicious choice will allow us to obtain consistency of Ξ^I as only I (the

²³ $\sqrt{N} E[\epsilon_t^i | p_t^I] \rightarrow \lambda$, so *a fortiori*, $E[\epsilon_t^i | p_t^I] \rightarrow 0$.

number of subjects) increases, keeping T (the number of periods in an experimental session) fixed, and, if possible, small.

We obtain risk tolerances from OLS projections of holdings onto $\Delta^{-1}(\mu - p_t^I)$. We use end-of-period holdings for all periods *except* period t (the period on which we run our GMM test). Let $\hat{\beta}_t^i$ denote our estimate of subject i 's risk tolerance. Define:

$$\hat{\beta}_{j,t}^i = \frac{\text{cov}([z_\tau^i]_j, [\Delta^{-1}(\mu - p_\tau^I)]_j)}{\text{var}([\Delta^{-1}(\mu - p_\tau^I)]_j)},$$

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

$$\hat{\beta}_t^i = \frac{1}{J} \sum_j \hat{\beta}_{j,t}^i.$$

The estimation error, $\hat{\beta}_t^i - \beta_t^i$, depends linearly on the perturbations ϵ_τ^i for all periods τ except $\tau = t$. To demonstrate this, consider the following:

$$\hat{\beta}_{j,t}^i - \beta_t^i = \frac{\text{cov}([\epsilon_\tau^i]_j, [\Delta^{-1}(\mu - p_\tau^I)]_j)}{\text{var}([\Delta^{-1}(\mu - p_\tau^I)]_j)}.$$

Therefore,

$$\hat{\beta}_t^i - \beta_t^i = \frac{1}{J} \sum_j \frac{\text{cov}([\epsilon_\tau^i]_j, [\Delta^{-1}(\mu - p_\tau^I)]_j)}{\text{var}([\Delta^{-1}(\mu - p_\tau^I)]_j)}.$$

The sample covariances in the last expression are linear in the perturbations $[\epsilon_\tau^i]_j$. It follows that the estimation error $\hat{\beta}_t^i - \beta_t^i$ is linear in the perturbations $[\epsilon_\tau^i]_j$.

Linearity implies that our estimator will be unbiased asymptotically because $E[\epsilon_t^i | p_t^I] \rightarrow 0$:²⁴ $E[\hat{\beta}_t^i - \beta_t^i | p_1^I, \dots, p_t^I] \rightarrow 0$.

Also, linearity, together with the assumed asymptotic conditional time series orthogonality of individual perturbations implies that the estimation

²⁴ $\sqrt{N}E[\epsilon_t^i | p_t^I] \rightarrow \lambda$, so *a fortiori*, $E[\epsilon_t^i | p_t^I] \rightarrow 0$.

error $\hat{\beta}_t^i - \beta_t^i$ is uncorrelated with ϵ_t^i :

$$\begin{aligned} & E\{[\epsilon_t^i]_k(\hat{\beta}_t^i - \beta_t^i)|p_1^I, \dots, p_T^I\} \\ &= \frac{1}{J} \sum_j \frac{\text{cov}(E\{[\epsilon_t^i]_k[\epsilon_\tau^i]_j|p_1^I, \dots, p_T^I\}, [\Delta^{-1}(\mu - p_\tau^I)]_j)}{\text{var}([\Delta^{-1}(\mu - p_\tau^I)]_j)} \\ &\rightarrow 0, \end{aligned}$$

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

The Impact of Estimation Error in Risk Tolerances

To demonstrate that the errors in estimating risk tolerances have no effect on Ξ^I asymptotically, first consider the leading factor in the definition of Ξ^I , namely, $(B_t^I)^2$. Since individual risk tolerances are estimated in an unbiased way,

$$\frac{1}{I} \sum_{i=1}^I \beta_t^i - \frac{1}{I} \sum_{i=1}^I \frac{1}{b_t^i} \rightarrow 0,$$

by the law of large numbers, so estimation error in the leading factor can be ignored asymptotically.

Ignoring the leading factor, consider next the second term in the formula for Ξ^I . (The argument for the third term is analogous and will not be presented.) Rewrite it in terms of the true risk tolerances plus estimation errors:

$$\begin{aligned} & \frac{1}{I} \sum_{i=1}^I (z_t^i - \bar{m}^I) \left([\hat{\beta}_t^i - \frac{1}{I} \sum_{\nu=1}^I \hat{\beta}_t^\nu] \Delta^{-1}(\mu - p_t^I) \right)^{\text{T}} \\ &= \frac{1}{I} \sum_{i=1}^I (z_t^i - \bar{m}^I) \left([\beta_t^i - \frac{1}{I} \sum_{\nu=1}^I \beta_t^\nu] \Delta^{-1}(\mu - p_t^I) \right)^{\text{T}} \\ & \quad + \frac{1}{I} \sum_{i=1}^I (z_t^i - \bar{m}^I) \left(\left[(\hat{\beta}_t^i - \beta_t^i) - \frac{1}{I} \sum_{\nu=1}^I (\hat{\beta}_t^\nu - \beta_t^\nu) \right] \Delta^{-1}(\mu - p_t^I) \right)^{\text{T}} \end{aligned}$$

Consider the deviations of portfolio choices from the grand mean in the second term of the last expression, $z_t^i - \bar{m}^I$, $i = 1, \dots, I$. These depend linearly

on the perturbations ϵ_t^i , $i = 1, \dots, I$. In the same term, the estimation errors, namely, $\hat{\beta}_t^i - \beta_t^i$ and $\hat{\beta}_t^\nu - \beta_t^\nu$ are asymptotically mean zero. They are also asymptotically uncorrelated with the perturbations ϵ_t^i , because they depend linearly on perturbations ϵ_τ^i for $\tau \neq t$, as demonstrated earlier. Clearly, the second term in the above expression is simply the sample covariance of linear transformations of perturbations ϵ_τ^i for $\tau \neq t$, on the one hand, and linear transformations of the perturbations ϵ_t^i , on the other hand. Asymptotically, this sample covariance converges to zero in expectation. Because perturbations ϵ_τ^i and ϵ_t^i are assumed independent across i , 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.

This leaves us only with the first term in curly brackets. The random behavior of the first term obviously does not depend on errors in the estimation of the risk tolerances. We have the desired result: asymptotically, estimation errors have no impact on Ξ^I .

Consistency of Ξ^I

Because their estimation errors have no effect asymptotically, we can write Ξ^I as a function of the true risk tolerances. This is what we did in (18). Convergence to W^I , and hence, A , is immediate. ■

References

- K. ARROW AND F. 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.
- J. BERK, "Necessary Conditions for the CAPM," *Journal of Economic Theory* 73 (1997), 245-257.
- P. BOSSAERTS AND C. PLOTT, "The CAPM in Thin Experimental Financial Markets," *Journal of Economic Dynamics and Control* 26 (2001), 1093-1112.
- P. BOSSAERTS AND C. PLOTT, "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," *Review of Finance* 8 (2004), 135-169.
- P. BOSSAERTS, C. PLOTT AND W. ZAME, "Discrepancies between Observed and Theoretical Choices in Financial Markets Experiments: Subjects' Mistakes or the Theory's Mistakes?" Caltech Working Paper (2002).
- G. CONNOR, "A Unified Beta Pricing Theory," *Journal of Economic Theory* 34 (1984), 13-31.
- J. DAVIS, E. FAMA AND K. FRENCH, "Characteristics, Covariances, and Average Returns: 1929 to 1997," *Journal of Finance* 55 (2002), 389-406.
- R. GALLANT, *Nonlinear Statistical Models*. New York: Wiley, (1987).
- J. GEANAKOPOLOS AND M. SHUBIK, "The Capital Asset Pricing Model as a General Equilibrium with Incomplete Markets," *Geneva Papers on Risk and Insurance* 15 (1990), 55-71.

- M. GIBBONS, S. ROSS AND J. SHANKEN, "A Test of the Efficiency of a Given Portfolio," *Econometrica* 57 (1989), 1121-1152.
- M. HARRISON AND D. KREPS, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory* 20 (1979), 381-408.
- S. HART, W. HILDENBRAND AND E. KOHLBERG, "Equilibrium Allocations as Distributions on the Commodity Space," *Journal of Mathematical Economics* 1 (1974), 159-166.
- W. HILDENBRAND, *Core and Equilibria of a Large Economy*. Princeton, N.J.: Princeton University Press (1974).
- C. HOLT AND S. LAURY, "Risk Aversion and Incentive Effects," *American Economic Review* 92 (2002), 1644-1655.
- K. JUDD AND S. GUU, "Asymptotic Methods for Asset Market Equilibrium Analysis," *Economic Theory* 18 (2001), 127-157.
- M. LEDOUX AND M. TALAGRAND, *Probability in Banach Spaces*. New York: Springer (1991).
- D. MCFADDEN, "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica* 57 (1989), 995-1026.
- ROBERT C. MERTON, "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41 (1973), 867-887.
- A. PAKES AND POLLARD, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica* 57 (1989) 1027-1057.
- R. RADNER, "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica* 40 (1972), 289-303.
- R. ROLL, "A Critique of the Asset Pricing Theory's Tests, Part I: On the Past and Potential Testability of the Theory," *Journal of Financial Economics* 4 (1977), 129-176.

- W. SHARPE, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance* 19 (1964), 425-442.
- G. R. SHORACK AND J.A. WELLNER, *Empirical Processes with Applications to Statistics*. Wiley, New York (1986).