Prices and Portfolio Choices in Financial Markets: Theory and Experiment*

Peter Bossaerts
California Institute of Technology
Centre for Economic Policy Research

Charles Plott
California Institute of Technology

William R. Zame
UCLA
California Institute of Technology

July 1, 2003
This version March 10, 2005

*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 John Simon Guggenheim Foundation, the Social and Information Sciences Laboratory at Caltech, and the UCLA Academic Senate Commitee on Research 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.

Prices and Portfolio Choices in Financial Markets: Theory and Experiment

Abstract

Most tests of asset pricing models address only the pricing predictions — perhaps because the portfolio choice predictions are obviously wrong. But how can pricing theory be right if the portfolio choice theory on which it rests is wrong? This paper suggests an answer: the assumptions about individual preferences that underly common asset-pricing models are wrong, but the deviations between the demands predicted by these models and the true demands have mean zero in the population, and hence wash out in prices.

The starting point for this work is a set of experimental markets in which risky and riskless assets are traded. This experimental setting offers an opportunity to study asset pricing in an environment in which crucial variables can be controlled or observed — in contrast to field environments, in which these variables cannot be controlled and frequently cannot be observed accurately. The experimental data exhibit the same puzzling characteristics as the historical data: asset prices are consistent with the price predictions of familiar theories (for instance, the market portfolio is nearly mean-variance efficient) but portfolio choices are wildly divergent from the portfolio choice predictions of the same theories (for instance, portfolio separation does not obtain). To explain the data, we build a structural model based on perturbations of individual demand functions (in the familiar style of much applied work). The central feature of this model is that the perturbations (i.e., the differences between individual demands and the demands predicted by mean-variance utility) have mean zero in the population. We develop an econometric test of the model which tests *both* prices and portfolio choices, and find that the empirical distribution of the test statistic is consistent with model predictions.

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

Keywords experimental finance, experimental asset markets, risk aversion

1 Introduction

Most asset pricing models predict both asset prices and portfolio choices. Forty years of econometric tests of such models present weak support for the pricing predictions, 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, these studies offer a puzzle: How can the price predictions of asset pricing models be right if the portfolio choice predictions of these same models are wrong?

In this paper we suggest a resolution to this puzzle. Our analysis has three parts. The first part presents data from experimental asset markets — an environment in which asset payoffs and the information available to investors can be controlled and prices and portfolio choices can be observed. The price data from these experiments are consistent with the pricing predictions of standard asset pricing models, including CAPM — in particular, the market portfolio is mean-variance efficient — but the portfolio choice data are not consistent with the portfolio choice predictions of the same models — in particular, investors do not hold the market portfolio. Indeed, the distance from actual portfolio choices to theoretical predictions of portfolio choices is uncorrelated with the distance from actual asset prices to theoretical predictions of asset prices. (These findings are perhaps all the more striking because subjects are not informed of the holdings of others or of the market portfolio, and hence cannot use standard asset-pricing models to predict prices.) The second part presents a simple theoretical model that is capable of explaining these data. Our model 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. In a large market, the idiosyncratic components average out across the population, and CAPM pricing prevails — but CAPM portfolio choice predictions do not. The third part takes this theory to the experimental data using novel econometric tests based on a cross-sectional version of GMM, and testing both prices and portfolio choices. We find that the empirical distribution of the test statistic is consistent with model predictions.

In our experimental markets, $\sim$ 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 probability that each state will occur, and of course they know their own holdings and their own attitudes toward wealth and risk. They also have access to the history of orders and trades. Subjects do not know the number of participants in any given experiment, nor the holdings of other participants, nor the market portfolio. We analyze our data in the context of the static Capital Asset Pricing Model (CAPM).[1] To do so, we follow the standard strategy, familiar from empirical studies of historical data, and use end-of-period prices and portfolio holdings, ignoring intra-period prices.[2] (Because our securities have only one-period lives, so that we can use liquidating dividends as security payoffs, while empirical tests of historical data usually take end-of-month prices as security payoffs, our experiments actually represent an environment that is closer to a static asset-pricing model than are typical field studies.) Our experimental data are consistent with the pricing predictions of CAPM: the market portfolio is (approximately) mean-variance efficient. On the other hand, our experimental data are *inconsistent* with the portfolio choice predictions of CAPM: individual investors do *not* hold the market portfolio. Indeed, individual portfolio holdings seem almost random. And, as we have said, the distance from actual portfolio choices to theoretical predictions of portfolio choices is uncorrelated with the distance from actual prices to theoretical predictions

[1]Other standard asset pricing models would yield similar implications.

[2]This 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.

of prices.

To explain these findings, we extend the standard CAPM to incorporate *unobserved heterogeneity*; we call the extended model CAPM+ϵ. Our approach is similar to that used in much applied work: we assume that demand functions of individual traders can be decomposed as sums of mean-variance components and idiosyncratic components (perturbations), and that these idiosyncratic components are drawn from a distribution that has mean zero. If these idiosyncratic components are independent and the population is large, the Law of Large Numbers implies that the perturbations (approximately) wash out in the aggregate. Hence CAPM+ϵ predicts the same equilibrium prices as does CAPM but is consistent with portfolio choices very different than those predicted by CAPM.

To test 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 drawn from a distribution that has mean zero. These tests are novel in that they link prices and portfolio choices, and also in other ways:

- In the usual models of choice with unobserved heterogeneity, the null hypothesis is that the idiosyncratic components of demand have mean zero and are orthogonal to prices. In our setting, however, it is the only the idiosyncratic components of demand *functions*, rather than *realized demands* that have mean zero. Because demands influence equilibrium prices, realized (equilibrium) demands 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 obtain a meaningful test, we need to estimate the unknown noncentrality parameter of the asymptotic distribution of our test statistic. To do this, we use our multiple experiments to generate multiple samples. Because samples from different periods within a single experiment are not independent, we construct empirical distribution functions using data *across experiments* rather than *within experiments*, and use standard Kolmogorov-Smirnov and Cramer-von Mises statistics to test whether the empirical distribution function could have been generated by a member of the family of non-central χ^2 distribution functions.

- 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.

Following this Introduction, Section 2 describes our experimental asset markets. Section 3 describes the data generated by our experiments, and discusses the relationship of these data to the standard CAPM. Section 4 describes our expanded theoretical model and Section 5 describes the econometric methodology and findings. (Technical details are relegated to Appendices.) Section 6 concludes.

2 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 only be held in non-negative amounts (cannot be sold short). Cash can only be held on 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. (In the experiments considered here, 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. In some experiments, states of nature for each period were drawn independently from the uniform distribution. Randomization was achieved by the use of a random number generator or by drawing balls from an urn, with drawn balls replaced. In the remaining experiments states were not drawn independently. Rather, balls marked with the state were drawn from an urn that initially contained 18 balls, 6 for each state, and drawn balls were not replaced. In each treatment, subjects were informed as to the procedure. Subjects are *not* informed of the endowments of others, or of the market portfolio (the social endowment of all assets), or the number of subjects, or whether these 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.) We were especially careful not to provide information about the market portfolio, so that subjects could not easily deduce the nature of aggregate risk — lest they attempt to use a standard model (such as CAPM) to *predict* prices, rather than to take observed prices as given. Keep in mind that neither general equilibrium theory nor asset pricing theory require that participants have any more information than is provided in these experiments. Indeed, much of the power of these theories comes precisely from the fact that agents know — 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. The state-dependent payoffs of assets (in francs) are recorded in the following table.

Table 1: Asset Payoffs

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

1 unit of Cash is 1 franc in each state of nature. The remaining parame-

ters for the various experiments are displayed in Table 2. Experiments are 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 (due to wealth effects, and possible effects of bankruptcy or the fear of bankruptcy), there is every reason to expect equilibrium prices to differ across periods in a given experiment. Note that, given the true probabilities, $\mathrm{cov}(A, B) < 0$; 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).[3] 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.

[3] Complete instructions and other details are available at http//eeps3.caltech.edu/market-011126; use anonymous login, ID 1, password a.

Table 2: Experimental Parameters

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

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

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

3 Experimental Data

In this Section, we summarize the data from our experimental markets, using a simple model to help organize it.

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.[4] 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.[5] 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 our experimental markets occurs throughout each period, no information is revealed during that time; information is only revealed 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.)

[4] A complete record of every transaction in every experiment is available at: http://www.hss.caltech.edu/∼pbs/BPZdata. Because the 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.

[5] The historical record provides little information about holdings.

3.1 CAPM

In our experiments, two risky assets and a riskless asset (Notes) are traded against Cash. However, Cash and Notes have the same payoffs. Simplifying slightly, we therefore treat Cash and Notes as perfect substitutes, hence redundant assets, and use a model with two risky assets and one riskless asset. (If Cash and Notes were exact perfect substitutes, then the price of Notes would be exactly 100 at the end of each trading period. As Table 3 shows, this is approximately the case in most periods in most experiments. That the price of Notes is not *exactly* 100 at the end of every period reflects the fact that all transactions must take place through Cash, so that there is a transaction value of Cash.)

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 $\operatorname{div} A$ for the state-dependent dividends of asset A, $\operatorname{div} A(s)$ for dividends in state s, and so forth. If $\theta = (\theta_A, \theta_B, \theta_N) \in \mathbb{R}^3$ is a *portfolio* of assets, we write

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

for the state-dependent dividends on the portfolio θ.

There are I investors. Investor i is characterized by an endowment portfolio $\omega^i = (\omega^i_A, \omega^i_B; \omega^i_N) \in \mathbb{R}^2_+ \times \mathbb{R}$ of risky and riskless assets, and a strictly concave, strictly monotone utility function $U^i : \mathbb{R}^3 \to \mathbb{R}$ defined over state-dependent terminal consumptions. (To be consistent with our experimental design, we allow consumption to be negative.) 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 $\operatorname{div} \theta^i$ subject to the budget constraint $q \cdot \theta^i \leq q \cdot \omega^i$.

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

- choices are budget feasible: for each i

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

- choices are budget optimal: for each i

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

- asset markets clear:

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

Because the stakes in our experiments are small (in comparison to current wealth and even more so in comparison to present value of lifetime wealth), it is natural to approximate true preferences of subjects by mean-variance preferences. That is, we assume investor i's utility function for state-dependent wealth x is the form

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

where expectations and variances are computed with respect to the true probabilities, and b^i is absolute risk aversion.[6] We assume throughout that risk aversion is sufficiently small that the utility functions U^i are strictly monotone in the range of feasible consumptions (or at least observed consumptions).

In our environment, mean-variance utilities imply the *Capital Asset Pricing Model* (*CAPM*). We summarize the relevant implications of CAPM here; see Appendix A for a more complete derivation. Write $M = \sum \omega^i$ for the *market portfolio of all assets*, $m = \sum(\omega^i_A, \omega^i_B)$ for the *market portfolio of risky assets* and $\overline{M} = M/I$, $\overline{m} = m/I$ for the respective *per capita portfolios*.

[6] An alternative would be to assume individual preferences are of the state-independent expected-utility family and use the quadratic approximation given by Taylor's theorem. At the scale of our experiment, the differences between the mean-variance approximation and the quadratic approximation are almost unobservable; we prefer the mean-variance approximation only for econometric convenience.

Write $\mu = (E(A), E(B))$ for the vector of expected dividends of risky assets and

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

for the covariance matrix of risky assets. 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:

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

$$z^i(\tilde{p}) = \frac{1}{b^i}\Delta^{-1}(\mu - \tilde{p}) \tag{2}$$

(The quantity $\left(\frac{1}{I}\sum\frac{1}{b^i}\right)^{-1}$ is frequently called the *market risk aversion* and $\left(\frac{1}{N}\sum\frac{1}{b^i}\right)$ 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 $E(\mathrm{div}\, m) - q \cdot m$ on the portfolio m is highest among all portfolios having variance no greater than $\mathrm{var}(\mathrm{div}\, m)$.[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.

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

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

3.2 Prices and Holdings

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.

To compare observed prices with the predictions of CAPM we need a convenient measure of the deviation of the market portfolio from mean-variance efficiency; Sharpe ratios provide such a convenient measure. 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 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. Figure 1 summarizes these deviations in all our experiments.

Because a typical experiment involves more than 30 subjects, displaying portfolio holdings of each subject in each experiment is impractical and un-

Table 3: End-Of-Period Transaction Prices

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

[a]Security.

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

[c]Notes.

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

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

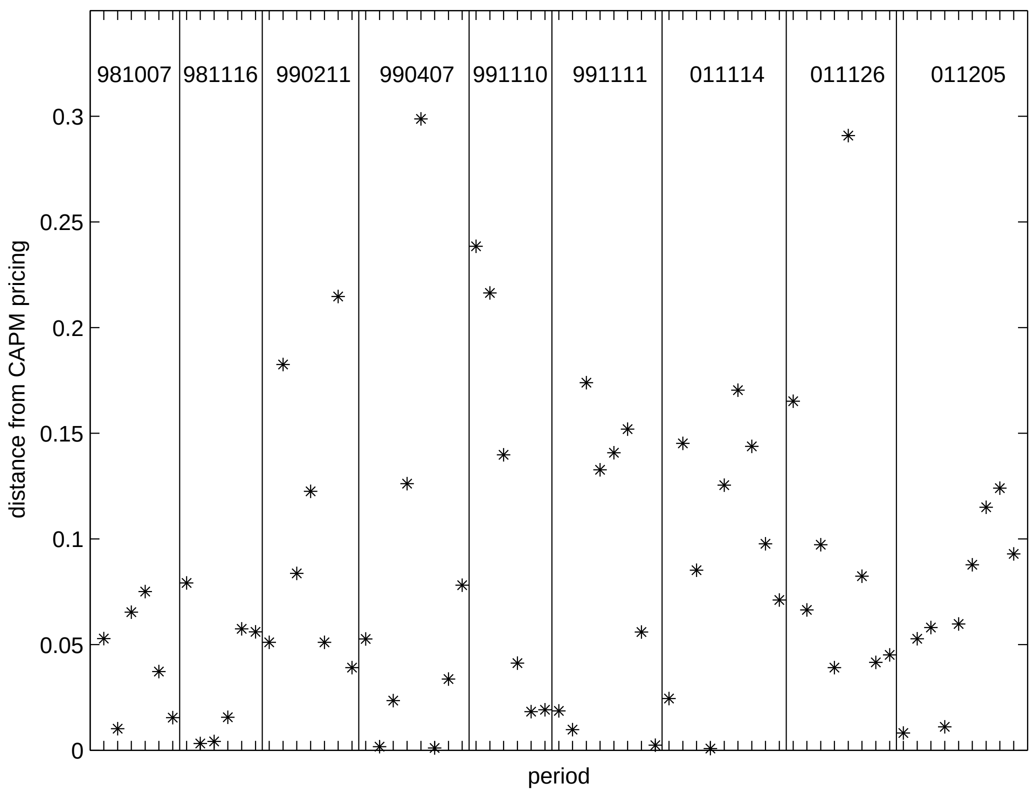

Figure 1: Plot of distance from CAPM pricing (measured as difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio) in all periods of all experiments

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:

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

Figure 2 displays mean absolute differences for each period in each experiment. As the reader can see, Portfolio Separation fails quite substantially. Indeed, the average deviations are roughly as large as they would be if investors chose portfolio weightings at random.[8]

As Figures 1 and 2 show, Mean Variance Efficiency seems confirmed in the experimental data while Portfolio Separation does not. An even more striking fact which is difficult to see in these tables, can be seen quite clearly in Figure 3. Each point (small circle) in Figure 3 represents a single period of a single experiment. The horizontal component of each point is the deviation (at the end-of-period prices) of the market portfolio from mean-variance efficiency; the vertical component of each point is the mean absolute deviation from portfolio separation. As can be seem very clearly in Figure 3, there is *no correlation* between the deviation from mean variance efficiency and the deviation from portfolio separation.

[8]To make the point 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 $E\subset[0,1]$. If this is the case then the mean absolute deviation will be only .24.

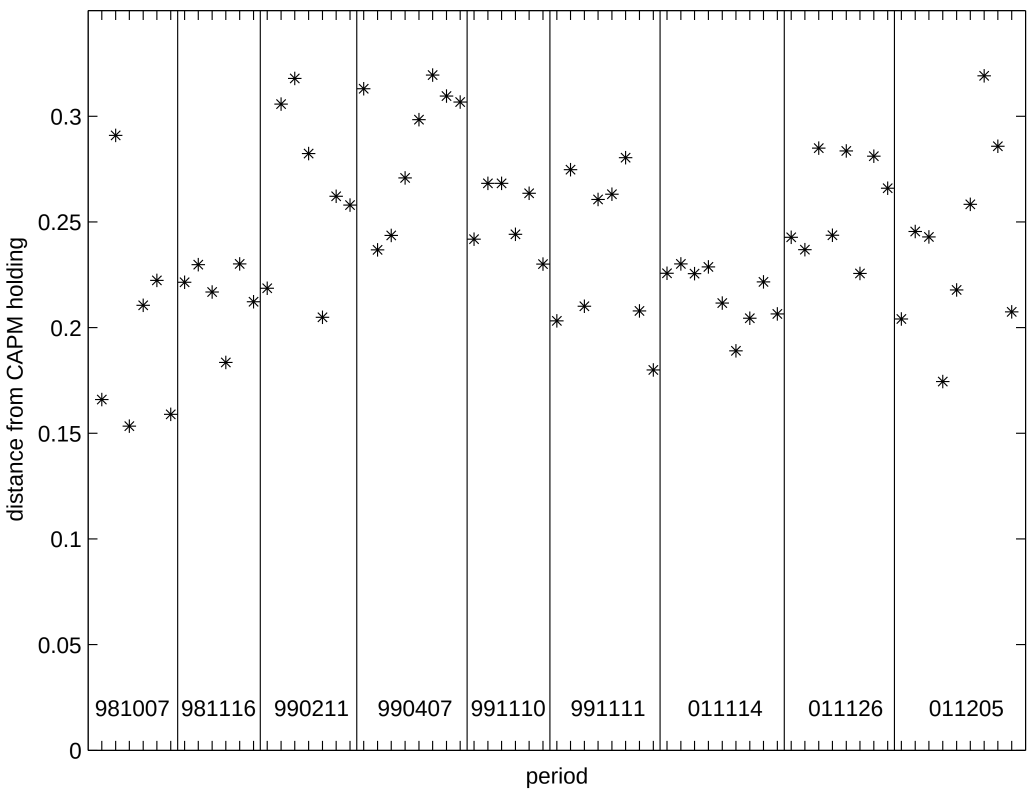

Figure 2: Plot of deviation from Portfolio Separation (measured as mean absolute difference between individual weighting on asset A and market weighting on asset A) in all periods of all experiments

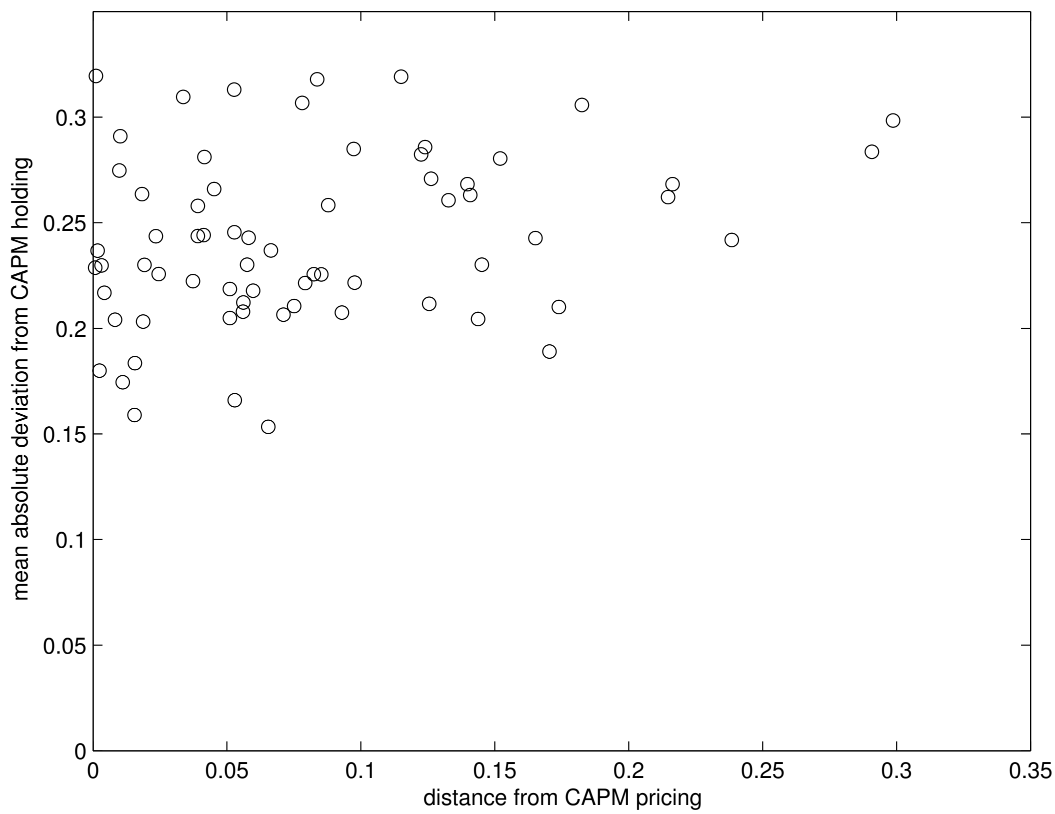

Figure 3: Plot of mean absolute deviations of subjects' end-of-period holdings from CAPM predictions against distances from CAPM pricing (absolute difference between market Sharpe ratio and maximal Sharpe ratio, based on last transaction prices), all periods in all experiments. There is no correlation between distance from CAPM pricing (x-axis) and violations of portfolio separation (y-axis).

4 CAPM+ϵ

In this Section we offer a model that is capable of explaining the data we have seen: it yields pricing predictions close to that of CAPM but is consistent with very different portfolio choices. Because the model differs from CAPM in that it adds perturbations (of demand functions), we refer to our model as CAPM+ϵ.

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 (true) demand structure in the marketplace. We implement this idea by viewing observed (true) demands as *perturbations* of hypothetical demands. In principle, these perturbations might represent some combination of subject errors (in computing and implementing optimal choices), market frictions and unobserved heterogeneity of true preferences. Because an adequate treatment of subject errors or market frictions would necessitate a fully stochastic model, which we are not prepared to offer, and because we have some evidence that subject errors and market frictions are not of most importance in our setting (see Bossaerts, Plott and Zame (2002)), we focus here on unobserved heterogeneity.

Because the approach we follow is quite intuitive, the following informal description is sufficient for our needs. Appendix B presents a careful and rigorous justification.

Consider an economy with I investors. Investor i has endowment portfolio ω^i and utility function U^i. Write $z^i(p)$ for i's demand for risky assets at prices p (which we assume to be single-valued). CAPM assumes that each U^i is a mean-variance utility function. Whether this is the case or not, however, we can always view U^i as a perturbation of a mean-variance utility function, and hence can write the true demand function z^i as a perturbation of a mean-variance demand function $\tilde{z}^i$:

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

In this way, we can view the true economy as a perturbation of a hypothetical economy in which the same number of investors have the same endowments but have mean-variance utilities. Write $D, \tilde{D}$ as the mean market excess demand functions in the true economy and in the hypothetical economy. By definition:

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

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 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 true economy and in the hypothetical economy coincide. More generally, if the mean perturbation $\frac{1}{I}\sum\epsilon^i$ is uniformly small, then equilibrium prices in the true economy and in the hypothetical economy nearly coincide. (This assertion requires some justification, which we provide in Appendix B.) Because the perturbations are drawn independently from a distribution with mean zero, a suitable version of the Strong Law of Large Numbers 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 hypothetical economy is mean-variance efficient at the equilibrium price $\tilde{p}$ of the hypothetical economy. Because the market portfolio of the true economy is the same as the market portfolio of the hypothetical economy we conclude that, if the number I of investors is large then, with high probability, the *market portfolio of the true economy will be approximately mean-variance efficient at the equilibrium prices of the true economy.* Of course, individual portfolio choices in the true economy need bear no obvious relationship to individual portfolio choices in the hypothetical economy; in particular, because the perturbations ε^i need not be small, approximate portfolio separation need not hold in the true economy.

Perhaps the most important feature of this model is that provides a mechanism leading to mean-variance efficiency of the market portfolio even though *no single investor* chooses a mean-variance optimal portfolio. (In the standard 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. Evidence for this suggestion can be found in Bossaerts and Plott (2002).

5 Structural Econometric Tests

In this Section, we construct a structural econometric test of CAPM+ϵ, and then apply this test to the data from our experiments. Our approach has a number of novel features that distinguish it from the usual approaches to the econometric analysis of field data:

- The usual approaches rely entirely on market prices. Our approach links market prices and individual holdings.
- In the usual approaches, randomness is viewed as the sampling error in estimation of the distribution of returns. In our approach, randomness is viewed as the deviations of observed choices from hypothetical 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. (We use this approach because we do not have long time series: our experiments are only 6-9 periods long, and it is simply impractical to conduct significantly longer experiments. However, we can exploit the fact that each experiment involves approximately 30-60 subjects, so that we have large cross-sections.)
- The usual approach is to test a theory on each sample separately, and then aggregate the results. Our approach is to construct samples as aggregates of periods in different 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. (We use this approach because we lack information about perturbation terms that would be necessary to uniquely identify the asymptotic distribution of our test statistic, and hence cannot test our theory on the data from a single period in a single experiment. Because outcomes across periods within an experimental session are not

independent — each period involves the same collection of subjects, and wealth effects in early periods may influence attitudes toward risk in later periods — our samples consist of outcomes in the same period across the various sessions: one sample consists of outcomes from the first period in each experiment, etc.)

An additional novel feature of our approach is in our estimation of coefficients of risk aversion for each individual in each period. To obtain these estimates, we use observed choices for that individual in *other* periods of the same experimental session. Because the number of periods in each session is small, our estimates cannot be accurate. However, we are able to use an estimation procedure with the property that estimation error tends to cancel out across subjects. This approach is reminiscent of one 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).

5.1 The Null Hypothesis

We focus 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; as in Section 4 and Appendix B, we view the characteristics of the subjects as drawn from a population with a given distribution. As before, we write μ for the vector of mean payoffs of the risky assets, Δ for the covariance matrix of payoffs of risky assets, and $\overline{m}^I$ for the per capital market portfolio of risky securities. The superscript makes explicit that $\overline{m}^I$ may vary with the number of subjects/investors I. This will facilitate econometric analysis: we are interested in asymptotic properties of our test statistic as $I \to \infty$. We subscript demand functions z_t^i to emphasize that they depend on the particular period; we do not subscript the quantities ω^i, μ, Δ, $\overline{m}^I$ because they do not depend on the particular period t. Because endowments 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 of investor having the same endowment as trader i and a mean-variance utility function with coefficient of risk aversion b_t^i. The difference between the true demand function z_t^i and the hypothetical demand function $\tilde{z}_t^i$ is the *perturbation* or *error*:

$$\epsilon_t^i = z_t^i - \tilde{z}_t^i \tag{3}$$

Let $\tilde{\mathcal{E}}_t$ be the economy populated by these mean-variance traders, and 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 economy $\tilde{\mathcal{E}}_t$. (We use the superscript I to emphasize that we have an economy with I investors.)

Assume that CAPM holds in the hypothetical economy $\tilde{\mathcal{E}}_t$ and write $\tilde{p}_t^I$ for the CAPM equilibrium prices (again, dependence on I is made explicit). At equilibrium, per capita demand must equal the per capita market portfolio so rewriting equation (2) in the present notation yields

$$\tilde{z}_t^i(\tilde{p}_t^I) = \frac{1}{b_t^i} \Delta^{-1} (\mu - \tilde{p}_t^I) \tag{4}$$

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

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

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:

$$\frac{1}{I} \sum z_t^i(p_t^I) = \overline{m}^I \tag{6}$$

Summing (4) and (5) over all investors i, combining with (6) and doing a little algebra yields the following relationship:

$$p_t^I \;=\; \tilde{p}_t^I + B_t^I \Delta \frac{1}{I} \sum_{i=1}^{I} \epsilon_t^i(p_t^I) \tag{7}$$

(Again, we use the superscript I to emphasize that we are considering an economy with I investors. Note that prices p_t^I appear on *both* sides of this equation so it is *not* a formula for equilibrium prices.)

It might seem natural to proceed as is common in applied work and take as null hypothesis the statement: The perturbations ϵ_t^i are mutually independent across i, given p_t^I, and

$$E[\epsilon_t^i | p_t^I] = 0. \tag{8}$$

This null hypothesis would 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 is the *wrong* null hypothesis: when I, the number of investors in the economy, is finite, market clearing condition 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, $E[\epsilon_t^i | p_t^I] \to 0$ as $I \to \infty$ — 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 that the conditional means of perturbations exhibit *Pitman drift*: for some λ,

$$\lim_{I \to \infty} \sqrt{I} \left(E[\epsilon_t^i | p_t^I] \right) = \lambda \tag{9}$$

Here we view the economy as a draw of I investors from a distribution of investor characteristics. So the expectation is taken over all investors in a particular draw, conditional on prices for that draw, and then over all draws. Under Pitman drift, the asymptotic distribution of the usual GMM statistic is non-central χ^2 with non-centrality parameter λ^2. Unfortunately, λ is unknown. As a result, CAPM+ϵ *cannot* be tested on a single sample (a single period). However, CAPM+ϵ *can* be tested based on the behavior of the GMM statistic across samples (periods), because the *form of its distribution* (non-central χ^2) is known.

5.2 Specifics of The GMM Statistic

Define

$$h_t^I(\beta) = \beta \frac{1}{I} \sum_{i=1}^{I} z_t^i(p_t^I) - \Delta^{-1}(\mu - p_t^I). \tag{10}$$

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:

$$\min_{\beta} [\sqrt{I} h_t^I(\beta)^{\mathrm{T}}] W^{-1} [\sqrt{I} h_t^I(\beta)], \tag{11}$$

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

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

$$\begin{aligned}
h_t^I(B_t^I) &= B_t^I \frac{1}{I} \sum_{i=1}^{I} z_t^i(p_t^I) - \Delta^{-1}(\mu - p_t^I) \\
&= B_t^I \frac{1}{I} \sum_{i=1}^{I} \left[z_t^i(p_t^I) - \frac{1}{b_t^i} \Delta^{-1}(\mu - p_t^I) \right] \\
&= B_t^I \frac{1}{I} \sum_{i=1}^{I} \epsilon_t^i(p_t^I) \qquad (12)
\end{aligned}$$

Hence, $E[h_t^I(B_t^I)|p_t^I] = B_t^I \frac{1}{I} \sum_{i=1}^{I} E[\epsilon_t^i|p_t^I] \to 0$, as asserted.) The solution of (11) therefore defines a *GMM estimator of the market risk aversion*: it generates the value β which makes the sample version of the expectation in (12) 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 (12) is almost surely strictly positive. Our criterion function [see (11)] will be strictly positive in large samples as well, because the sample version of (12) 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*.

5.3 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 (10) is zero. (To see this, note that $\frac{1}{I}\sum_{i=1}^{I} z_t^i$ in (10) is the mean demand for risky securities; at equilibrium prices, equals the market portfolio. Hence, if $h_t^I = 0$, the first-order conditions for mean-variance optimality are satisfied.) In particular, 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

$$\sqrt{I}h_t^I(B_t^I) = B_t\sqrt{I}\frac{1}{I}\sum_{i=1}^{I}\epsilon_t^i \tag{13}$$

(see equation (12)). 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.

5.3.1 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 $I \to \infty$) 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 concistency even for fixed T. Appendix C discusses our procedure in more detail.

5.3.2 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.

We cannot readily aggregate the results over all periods, because the GMM statistics across periods within an experiment are not independent (because of wealth effects and because we estimate individual risk tolerances from choices in other periods, among other reasons). Fixing a period, however, the GMM statistics can safely be assumed to be independent across experiments. (The subject populations of different experiments are disjoint.) Thus, we test whether the empirical distribution of GMM statistics is non-central χ^2 *for a given period.*

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.[9]

There are two reasons why our test should be considered to be powerful.

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.

[9] An 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.

5.4 Test Results

Table 4 reports Kolmogorov-Smirnov (KS) and Cramer-von Mises (CvM) tests of whether the empirical distribution functions of our GMM statistics for a fixed period across experiments is non-central χ^2 with the best-fitting non-centrality parameters (11.6 for for KS, 10.0 for CvM).[10] All statistics are corrected for small-sample biases as suggested in Shorack and Wellner (1986) [p. 239]; p values are obtained from the same source.[11]

At the 1% level, both KS and CvM goodness-of-fit tests reject only in period 2; both tests fail to reject in other periods. At the 5% level, KS rejects only in periods 1, 2, 5 while CvM rejects only in periods 1, 2; both fail to reject in other periods. The data therefore appear to support CAPM+ϵ.

Lest the reader find the p values in Table 4 smaller than one might hope, it may be useful to keep in mind that the p values derived in econometric tests of models on the basis of field data are usually *much* smaller (despite the fact that our tests are more stringent, in the sense that 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 $p = .005$ 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."

Figure 4 depicts the empirical distribution of the logarithm of the GMM

[10]Best 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

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

[11]Since critical values for the Cramer-von Mises statistic are known only for specific p values, we report the range in which a p value fall.

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

Period Number[e]	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	1.50	$0.05 > p > 0.025$	0.31	$0.15 > p > 0.10$
6	9	1.06	$p > 0.15$	0.39	$0.10 > p > 0.05$
7	6	0.96	$p > 0.15$	0.11	$p > 0.15$
8	4	1.26	$0.10 > p > 0.05$	0.42	$0.10 > p > 0.05$

[a]Kolmogorov-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].

[b]Based on Table 1 on p. 239 of Shorack and Wellner (1986).

[c]Cramer-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].

[d]Based on Table 1 on p. 239 of Shorack and Wellner (1986).

[e]Period 9 is not listed because of insufficient sample size.

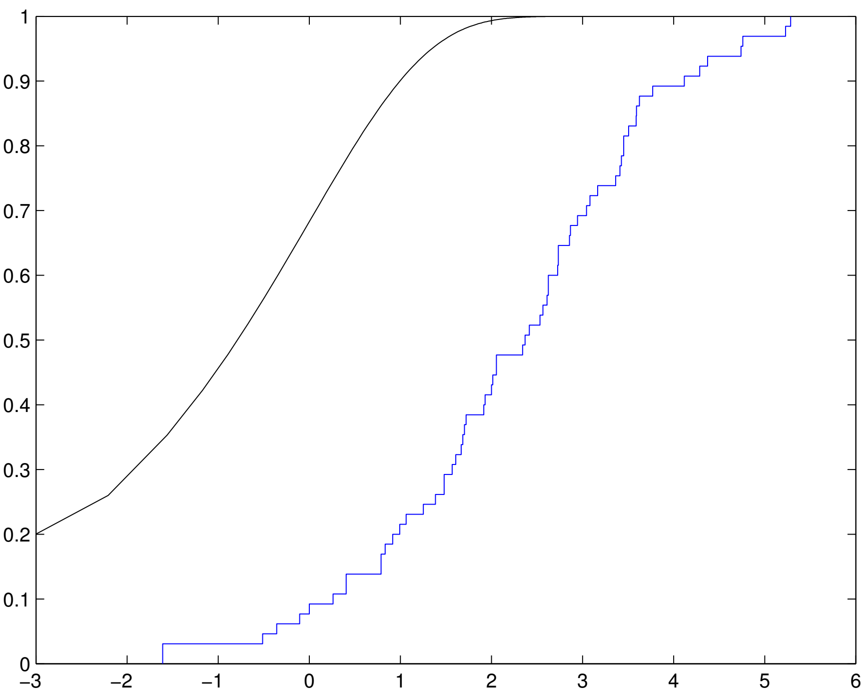

Figure 4: Empirical distribution of the GMM statistic, all periods in all experiments (jagged line), against a *central* χ^2 distribution. A noncentral χ^2 distribution provides a better fit, consistent with the small-sample biases expected if CAPM+ϵ is correct.

statistic across all periods in our experiments. For comparison, the smooth line represents the distribution of the logarithm of a *central* χ^2-distributed random variable. The empirical distribution (jagged line) appears to be a horizontal translation of the latter. This suggests that the GMM statistics are drawn from a non-central χ^2 distribution, which is confirmed in Table 4.

To gain further perspective, Tables 5 and 6 show GMM statistics and estimates of the harmonic mean risk aversion for all periods across all experiments. (We show experiments where states were drawn independently and experiments where states were drawn without replacement only because the combined table would be too large to display legibly on a single page.) p-values are provided based on the *central* χ^2 distribution (which ignores the correlation between prices and perturbations inherent to CAPM+ϵ). Put differently: p values in Table 5 are computed under the assumption that the non-centrality parameter $\lambda = 0$.

Note that the estimates β^I of the market mean risk aversion B_t^I are of the same order of magnitude across experiments, and are almost uniformly positive and significant: risk neutrality is rejected. This confirms our interpretation of the relation of prices to expected payoffs as reflecting "significant" risk premia.

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^2_1	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^2_1	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^2_1	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^2_1	7.5	0.7	13.8	116.7	†[b]	2.5	13.6	†	
	p level for $\lambda = 0$	.01	.39	.00	.00	-	.62	.00	-	
	β^I (*10^{-3})	0.5	0.5	0.3	-0.3	2.8	0.3	0.2	0.9	
	s.e. (*10^{-3})	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.1	
991110	χ^2_1	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^2_1	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	37.5	2.3	2.2	5.4	17.6	15.4	10.7	21.8	4.4
	p level for $\lambda = 0$	.00	.13	.14	.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	10.4	17.4	13.8	13.0	15.3	7.8	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	0.0	0.0

6 Conclusion

This paper has argued that findings from experimental financial markets provide significant insight into asset-pricing theory. Specifically, these findings demonstrate that deviations of observed individual demands from hypothetical mean-variance demands are idiosyncratic, and hence have little effect on prices but may have an enormous effect on portfolio holdings. These findings suggest that it makes sense to test asset pricing models that rely on strong portfolio separation results, such as CAPM and its multi-factor extensions, even in the absence of convincing evidence for such portfolio separation.

As a by-product of our structural tests, we provide evidence that the standard model of choice under unobserved heterogeneity in applied economics needs to be re-considered. In smaller markets, unexplained heterogeneity in demands (usually a key determinant of the unexplained portion of observed choices) need *not* be orthogonal to prices, and such non-orthogonality may have significant effect on the econometric analysis.

The econometric procedure introduced here explicitly links prices to allocations, unlike in standard field tests of asset pricing theory. With representative data on portfolio choices, this procedure could be applied to field data as well.

Appendix A: CAPM

To derive the conclusions of CAPM in our setting in which sort sales of risky assets are not permitted, we begin by analyzing the setting in which arbitrary short sales are permitted. Write $\widehat{z}^i(p)$ for investor i's demand for risky assets when the price of risky assets is p. (We suppress demand for the riskless asset because it is determined by budget balance.) 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

$$\widehat{z}^i(p) = \frac{1}{b^i}\Delta^{-1}(\mu - p) \tag{14}$$

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 $\widehat{p}$:

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

In our setting, short sales of risky assets are not permitted, and demand functions are *not* given by the equation (14). 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 $\widehat{z}^j(p^*)$

for investor j. However, formula (14) guarantees that if $\widehat{z}^j(p^*)$ were positive for some investor j then $\widehat{z}^i(p^*)$ would be positive for every investor i, whence $z^i(p^*)$ would coincide with $\widehat{z}^i(p^*)$ for every investor i. Because $p^* \neq \widehat{p}$, 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:

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

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

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

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

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

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

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

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

Appendix B: CAPM+ϵ

In this Appendix we give a formal and rigorous presentation of the idea of the true economy as draw from a distribution of individual characteristics and as a perturbation of a mean-variance economy. Although the ideas are very simple, and the conclusions both intuitive and expected, the details require a little care.

We work throughout in the setting of Section 4, 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 $End \subset \mathbb{R}^2_+ \times \mathbb{R}$, $\mathcal{P} \subset \mathbb{R}^3_{++}$, $[0, \overline{w}] \subset \mathbb{R}_+$. An investor is characterized by an endowment $\omega \in End$ of riskless and risky assets and by a continuous demand function

$$Z : \mathcal{P} \times [0, \overline{w}] \to \mathbb{R}^2_+ \times \mathbb{R}$$

for risky and riskless assets as a function of wealth $w \in [0, \overline{w}] \subset \mathbb{R}_+$ and prices for risky assets $p \in \mathcal{P} \subset \mathbb{R}^2_{++}$. (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:

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

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 $(\omega^i, Z^i) \in \operatorname{supp} \tau \subset End \times \mathcal{D}$. We identify the set of such draws with $(\operatorname{supp} \tau)^I$, 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 $(\operatorname{supp} \tau)^I$ 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 $(p,1)\cdot\theta \le (p,1)\cdot\omega$ for every $p \in \mathcal{P}$. We say $\sigma \in \mathcal{M}(End \times \mathcal{D})$ is a *mean-variance distribution* if for each $(\omega, Z) \in \operatorname{supp}\sigma$ there is a coefficient of risk aversion $b(\omega, Z) > 0$ such that the mean-variance utility function

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

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 $\overline{m}_\sigma = \int \omega d\sigma$ for the *per capita market portfolio.* We say the mean-variance distribution σ is a *CAPM distribution* if the price

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

belongs to the interior $\operatorname{int}\mathcal{P}$ 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 portolio $\overline{m}$ 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 $\overline{m}$ (portfolio separation).

Mean Zero Perturbations

Write $\pi_E : E\times\mathcal{D} \to E$ 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 $f : \operatorname{supp}\tau \to \mathcal{D}$ 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 $\Gamma_I \subset (supp\,\tau)^I$ 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 $D_\nu : \mathcal{P} \to \mathbb{R}^2_+ \times \mathbb{R}$ be the market demand function for assets:

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

and let $\overline{M}_\nu$ be the per capita market portfolio

$$\overline{M}_\nu \;=\; \int \omega d\nu$$

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

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

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

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

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

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

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 $H : \mathcal{P} \to \mathbb{R}^2_+ \times \mathbb{R}$ is any continuous function for which

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

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

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

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 $\Phi : \mathcal{P} \to \mathbb{R}^3$ (see Ledoux and Talagrand (1991) for the appropriate Banach space version) and in $\mathbb{R}^3$ implies that there is an index I_0 such that if $I > I_0$ then

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

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

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

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:

$$E\{[\epsilon_t^i]_k[\epsilon_\tau^i]_j|p_1^N, ..., p_t^I\} \to 0, \tag{15}$$

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

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

$$\overline{m}^I = \frac{1}{I}\sum_{i=1}^{I} z_t^i. \tag{16}$$

Also, define:

$$W^I = IE[h_t^I(B_t^I)h_t^I(B_t^I)^{\mathrm{T}}|p_t^I].$$

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

$$W^I \;=\; IE[h_t^I(B_t^I)h_t^I(B_t^I)^{\mathrm{T}}|p_t^I]$$

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

where the last equality follows from

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

Now consider:

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

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

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

which converges to zero, because $E[\epsilon_t^i | p_t^I] \to 0$.[12] The same applies to the last term in the above expression.

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

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

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

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, $\mathrm{cov}(z_t^i)$:

$$\mathrm{cov}(z_t^i) = \frac{1}{I}\sum_{i=1}^{I}\left(z_t^i - \overline{m}^I\right)\left(z_t^i - \overline{m}^I\right)^{\mathrm{T}}.$$

Then let

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

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

$$-\frac{1}{2}(B_t^I)^2\frac{1}{I}\sum_{i=1}^{I}\left([\beta_t^i-\frac{1}{I}\sum_{\nu=1}^{I}\beta_t^\nu]\Delta^{-1}(\mu-p_t^I)\right)\left(z_t^i-\overline{m}^I\right)^{\mathrm{T}}. \quad (17)$$

Estimating Risk Tolerances

In oder 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 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:

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

where cov and var denote the sample covariance and variance, respectively, over τ in $1,...,T$ with $\tau\neq t$. Also, $j=1,...,J$, with J denoting the number of risky securities (length of the vector z_τ^i). $[y]_j$ denotes the jth 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:

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

Therefore,

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

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 $E[\epsilon_t^i | p_t^I] \to 0$:[13] $E[\hat{\beta}_t^i - \beta_t^i | p_1^N, ..., p_t^I] \to 0$.

Also, linearity, together with the assumed asymptotic conditional time series orthogonality of individual perturbations implies that the estimation error $\hat{\beta}_t^i - \beta_t^i$ is uncorrelated with ϵ_t^i:

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

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

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,

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

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:

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

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

$$+\frac{1}{I}\sum_{i=1}^{I}\left(z_t^i-\overline{m}^I\right)\left(\left[(\hat{\beta}_t^i-\beta_t^i)-\frac{1}{I}\sum_{\nu=1}^{I}(\hat{\beta}_t^\nu-\beta_t^\nu)\right]\Delta^{-1}(\mu-p_t^I)\right)^{\mathrm{T}}$$

Consider the deviations of portfolio choices from the grand mean in the second term of the last expression, $z_t^i-\overline{m}^I$, $i=1,...,I$. These depend linearly on the perturbations ϵ_t^i, $i=1,...,I$. 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 are 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 (17). 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. Geanakoplos 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. Harrison and D. Kreps, “Martingales and Arbitrage in Multiperiod Securities Markets,” *Journal of Economic Theory* (1979) 20, 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.

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.

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.

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).