Testing CAPM in Real Markets: Implications from Experiments

Peter Bossaerts[‡]

This version: May 2003; First version: October 2002

[‡]California Institute of Technology and CEPR. Address: Caltech m/c 228-77, Pasadena, CA 91125. Phone: 626-395-4028. Fax: 626-405-9841. Email: pbs@rioja.caltech.edu.

Testing CAPM in Real Markets:
Implications from Experiments

Abstract

Tests of the CAPM, the prototype model of equilibrium in financial markets, are usually based on returns computed from end-of-month closing prices. It is reasonable to doubt that these prices always reflect markets that are at equilibrium, thus raising the question whether and how inference is biased. Rather than exploring this issue using one of the many theoretical (but empirically unverified) equilibration models that have been suggested in the literature, this paper starts from an empirical analysis of equilibration in experimental competitive markets. Price and allocation dynamics reveal systematic patterns which can be explained in terms of an equilibration model that builds on a simple and plausible behavioral premise, namely, *local optimization.* To get a comprehensive understanding of the dynamics of prices and allocations in disequilibrium, an extreme version of the equilibration model is studied, whereby prices move much faster than trades, so that all transactions occur at local equilibrium. A portfolio is identified that remains on the mean-variance efficient frontier throughout the equilibration process, namely, the risk-aversion-weighted endowment (RAWE) portfolio. The RAWE and the market portfolio are closely related (the two eventually converge), so that the latter may stay close to the frontier as well. The experimental data confirm this as well as another result: the off-equilibrium equity premium may vastly over-estimate investors' risk aversion. Violations of portfolio separation (observed both in experiments and field data) are a natural consequence when equilibration dynamics stop short of (global) equilibrium because further gains from trade are insufficient. Explicit formulation of off-equilibrium allocation dynamics reveals in what direction the RAWE portfolio differs from the market portfolio: if payoff covariances are non-negative and everyone starts out with the market portfolio, then RAWE overweighs securities with low payoff variance. This implies, among other things, that the market portfolio will be out-performed when it is combined with a portfolio that is long in low-volatility and short in high-volatility securities.

Testing CAPM in Real Markets:
Implications from Experiments†

1 Introduction

In real financial markets, prices may move very fast, but allocations clearly don't. Most trades are small relative to total daily volume and daily volume is small relative to total capitalization. So, one wonders whether real financial markets clear instantaneously, as posited in the theory. Because of the frequency with which real markets are shocked by news, they may in fact never reach equilibrium (a view that defines the Neo-Austrian school of economic thought, building on F.A. Hayek's critique of neoclassical economics[1]). Why, then, would we care to bring equilibrium pricing models to data from real markets? Perhaps the question is of little relevance, because one could claim that the observation error caused by disequilibrium is small relative to averages and volatilities of returns over the typical observation interval (one month). But is it?

How do we know? We cannot just assume that observed prices are equilibrium prices plus an observation error and derive the implications of the ensuing errors-in-variables on standard tests of asset pricing models. Such an approach would be completely pointless, because implications would arbitrarily change with the ad-hoc assumptions on the error term. The alternative would be to use one of the many equilibration models that have been suggested in the literature (such as walrasian tatonnement). Unfortunately, none has a solid empirical foundation. They were actually constructed as tools to investigate stability of competitive equilibrium,[2] an exercise that up to today seems to only have engaged theorists.[3]

In this paper, we let the data go first. The data do not come from field markets, but from financial markets experiments, since only in the latter context can one unambiguously determine whether markets have equilibrated (prices and allocations can be observed and equilibrium prices and allocations can be computed). Certain phenomena are recorded that appear to be remarkably robust across parameterizations. A behavioral model is then conjectured, on which a theory of equilibration is built that is capable of explaining the observed phenomena. We subsequently derive the model's implications for tests of asset pricing theory on disequilibrium prices. This is done in an extreme version of the model, whereby prices move much faster than trades, so that all transactions occur at local equilibria (to be defined shortly).

†Financial support from the U.S. National Science Foundation (grant SES-0079374), the R.G. Jenkins Family Fund, and the William D. Hacker Chair in Economics and Management are gratefully acknowledged. The paper benefited from discussions with John Ledyard, Charles Plott, Bill Zame and Luigi Zingales. Errors are solely the responsibility of the author.

[1]See, e.g., Benink and Bossaerts (2001).

[2]See, e.g., Arrow and Hahn (1971).

[3]A major exception is the experimental work by Charles Plott. See Plott (2000), Anderson, e.a. (2003).

The experiments were designed to study the Capital Asset Pricing Model (CAPM). CAPM will be the focus of analysis, even if it is "out of fashion," being displaced by multi-factor models, a situation that should be attributed to the findings reported in Fama and French (1992). We will eventually be able to say something about multifactor models as well, however. In addition, we will draw conclusions about allocations and resolve a paradoxical situation which is not unique to CAPM, but shared with multi-factor models, namely: why would one be interested in verifying the pricing restrictions in these models if their allocation predictions are vastly at odds with the data?

We will focus on experiments designed to study the certainty-equivalent version of the CAPM, where quadratic payoff functions replace risk. Absent information about subjects' actual risk attitudes, it is preferable to directly induce the preferences at the core of CAPM (mean-variance utility), instead of exposing subjects to risk and merely assuming that utility is of the mean-variance type.[4] Even if subjects' preferences were mean-variance when faced with risk, we still would need to know subjects' risk tolerances in order to compute theoretical equilibrium prices and allocations. In the certainty-equivalent case, equilibrium prices and allocations can be computed, so that we will know unambiguously when the market is in equilibrium. This way, we can study when and how the market moves to equilibrium.

Our experimental data reveal that deviations from equilibrium pricing can be substantial. It is not uncommon to observe transactions at prices as much as 10% away from equilibrium levels. This should make it possible to identify salient features of the equilibration process, which turn out to be as follows. (i) Risk aversion (more precisely, the parameter that would measure risk aversion if there were risk instead of nonlinear payoffs) implicit in off-equilibrium prices is higher than at equilibrium, i.e., the equity premium is higher; implied risk aversion decreases gradually as time passes. (ii) Transaction prices of a security correlate systematically with excess demand of *all* securities, not only its own excess demand. (iii) The market portfolio is close to mean-variance optimal throughout the equilibration process, even at prices that differ significantly from equilibrium levels. In terms of allocations, the noteworthy phenomenon is that (iv) the portfolio separation on which CAPM relies fails even when prices are close to equilibrium levels.[5]

In Asparouhova, Bossaerts and Plott (2003), (ii) was documented in the context of markets experiments with risky assets, but no explanation was provided. Here, we introduce a model that not only implies (ii) but the other phenomena as well. In addition, it suggests new things to look at, which we shall do. Our model starts from the

[4]As a matter of fact, the assumption of mean-variance preferences is not too far off the mark. See Bossaerts, Plott and Zame (2002) for evidence. Mean-variance preferences show up in a laboratory context perhaps because risk, while not negligible, is nevertheless small, so that mean-variance optimization approximates subjects' actual desires fairly well.

[5]The latter finding is also pronounced when there is uncertainty. Bossaerts, Plott and Zame (2002) explores how one can get CAPM pricing and significant violations of portfolio separation in a static general equilibrium setting. The explanation rests on the introduction of perturbation terms, i.e., random deviations between observed and theoretical demands. The explanation is silent, however, about the origin of the perturbations. This paper should be interpreted as providing a specific explanation: the perturbations originate in optimal off-equilibrium trades. This will become clear later on.

behavioral premise that the (usually small) orders and trades in the marketplace do not reflect the demands and supplies of neoclassical economic theory. Instead, they correspond to marginal improvements in agents' portfolios, as if following a hill-climbing search subject to a quadratic adjustment cost. What we are conjecturing here is a simple model of behavior, namely, local optimization. We then derive a model of price and allocation dynamics that appears to fit the data. An open question is why subjects would engage in local optimization in competitive markets, instead of, e.g., submitting orders that are a fraction of their global (excess) demand. The following is a possibility. In our experiments, subjects face either a new situation or one that may have changed in ambiguous ways. Subjects are given the opportunity to re-trade, but do not know at which prices. The ambiguity of the situation may make it impossible for them to assign definite probabilities to all possible future trading opportunities. Is local optimization the right response?

Under local optimization, agents submit orders in the direction of maximum gain in utility, subject to a quadratic adjustment cost, and constrained by the prevailing prices in the market place. Order imbalances lead to price changes. The model is developed here in the context of the CAPM, and fine-tuned to the situation when prices move much faster than trades, so that all trades take place at a local equilibrium – to be understood as equilibrium for local demands and supplies, not the (globally optimal) demands and supplies of neoclassical economic theory. This generates interesting insights into, in particular, the structure of pricing and holdings off-equilibrium. That is, the model allows us to address the question we posed at the beginning.

Local optimization is also at the core of the equilibration model studied in Smale (1976), who demonstrates that the resulting process has nice asymptotic properties: unless it gets stuck, the limit points are Pareto optimal allocations. Additional experimental evidence in favor of local optimization as a behavioral foundation for equilibration dynamics is reported in Bossaerts, Plott and Zame (2003).

Heterogeneity in preferences is assumed throughout the entire analysis. Because of this, however, two phenomena (the cross-security correlation between price changes and excess demands, and the off-equilibrium mean-variance efficiency of the market portfolio) emerge only because the variables at hand (excess demands and the market portfolio) proxy for different, more fundamental variables (risk-aversion-weighted individual excess demands and the risk-aversion-weighted endowment portfolio, respectively). The theory essentially suggests that we looked at the wrong variables when exploring dynamics in experimental financial markets, and that the only reason we observed something systematic is because the proxies are highly correlated with the fundamental variables. We can go back to the data and verify whether this is indeed the case. If so, it is fair to conclude that we were lucky to find anything.

In terms of *pricing*, the most important result is the identification of a portfolio that remains on the mean-variance frontier throughout the equilibration process, namely, the risk-aversion-weighted endowment (RAWE) portfolio, i.e., average of each investor's holdings, weighted by relative risk aversion (portfolios of more risk averse investors are weighted more heavily). If, for instance, more risk averse subjects hold more value stocks, then the RAWE portfolio will overweigh value stocks (and underweigh growth stocks). All securities can be priced using the RAWE portfolio, even off-equilibrium (off global equilibrium, that is). Since RAWE can (trivially) be re-written as the market portfolio plus a zero-investment adjustment portfolio, securities can be priced using a two-factor pricing model with the market as one of the factors and the adjustment portfolio as the other one. In the example, the adjustment portfolio is long value stocks and short growth stocks.

In terms of *allocation dynamics*, the model generically predicts violations of portfolio separation even if everyone starts out with the same holdings. Investors are posited to adjust their portfolios into the direction of maximum local gain in utility (they are subject to quadratic adjustment costs), and not into the direction of globally optimal holdings. These directions may differ substantially. Take, for instance, the case where all investors start out with the market portfolio. Their globally optimal portfolios of risky securities, and hence, their demands (supplies) for risky securities, are perfectly proportional. Transactions occur whenever risk aversion differs, and if all realize a proportion of their excess demand (which, of course, can occur only at global equilibrium prices), they end up holding the same portfolio of risky securities, namely, the market portfolio. With local optimization, however, more risk averse investors are foremost interested in unloading the most risky securities in their portfolio in return for risk free securities, while the less risk averse investors sending in orders to buy the most risky securities against risk free securities, because such trade provide maximum marginal utility. Note that already after one transaction round (at local equilibrium prices), portfolio holdings differ across investors: more risk averse investors underweigh the most risky stock, less risk averse investors overweigh those. The extent to which volume is skewed towards the most risky stock depends on the correlation between the payoffs.

It is easy to see what this leads to if adjustment stops short of global equilibrium. Since effects from diversification are only second-order (as opposed to the first-order effects of establishing the right exposure to risk), adjustment may halt for lack of sufficient further utility gains. This occurs when investors have bought (or sold) the right number of risk free securities and additional gains only come from diversification. At that point, investors will hold vastly different portfolios and portfolio separation appears to fail dramatically. Yet, investors' action are perfectly consistent with the CAPM world. Failure of portfolio separation contrasts markedly with the success of CAPM pricing: by the time the equilibration process halts, the market portfolio is close to mean-variance optimal.

Rather than presenting the formal model and its implications first and then discuss the (experimental) evidence in its favor, the analysis will be presented as a dialogue between the theory and the data (not unlike the actual chronology of the discovery process that led to this paper). Thus, the next section discusses the phenomena in the experimental data. Section 3 conjectures local optimization, derives a model of equilibration and documents how this model explains the observed empirical phenomena. Section 4 discusses further implications for price behavior, and Section 5 discusses implications for the dynamics of portfolio holdings; in both cases, we discuss the extent to which the additional implications are borne out in the data. Implications for tests of CAPM on field data are derived in

Section 6. Section 7 concludes.

2 Experimental Background

2.1 Nature of the Experiments

Each experiment consists of a number of independent replications, referred to as *periods*, of the same situation. At the beginning of a period, subjects are endowed with a number of each of 3 securities, referred to as A, B and *Notes*, as well as cash. Subsequently, markets in each of the securities are opened, and subjects can submit orders and trade as they like, for a pre-set amount of time. The trading interface is a fully electronic (web-based) version of a continuous electronic open book. After markets are closed, subjects are paid depending on their final holdings of the securities, minus a fixed, pre-determined loan payment. The payment is nonlinear in the holdings of A and B, and linear in the holding of Notes, as detailed below. After payment, securities are taken away and a new period starts. Subjects keep the payments accumulated over the periods.

In contrast to traditional experiments in economics, subjects are not present in a centralized laboratory equipped with computer terminals, but access the trading platform over the internet. Communication takes place by email, phone and announcements through the main experiment web page.[6] That is, the experiment takes place in cyberspace, as opposed to a single physical location. This allows for much larger experiments, as the number of participants is only limited by server capacity. The larger scale ensures that a trading environment is created that approximates the perfect competition of the theory: bid–ask spreads are reduced to a minimum (one tick) and individuals have only marginal influence on pricing.

End-of-period payments are determined using payoff functions that are quadratic in the holdings of A and B and linear in the Notes. The chosen payoff function would be the utility function of an individual with mean-variance preferences if A and B were risky securities with a certain expected payoff and covariance structure, and Notes were risk free. Such mean-variance preferences lead to CAPM in equilibrium, as explained below. Let α be a scalar (used to denote the payoff on a single Note), μ a 2×1 vector of constants (these would be the expected payoffs if A and B were risky securities) and Ω a 2×2 symmetric, positive definite matrix (which would be the covariance matrix of the payoffs on A and B if there were uncertainty). Subject n, when holding h units of the Notes and the vector x of A and B will receive the payoff

$$u_n(h, x) = \alpha h + [x \cdot \mu] - \frac{b_n}{2}[x \cdot \Omega x] \tag{1}$$

b_n is a positive scalar (reflecting n's risk aversion under conditions of risk). Note that only b_n is subject-dependent.

[6] The interested reader can browse http://eeps3.caltech.edu/market-020528. This web site provides the instructions for a typical experiment (the 28 May 02 experiment), trading interface, and announcements. To log in as observer, use ID=1, password=a.

In the experiments, $\alpha = 100$,

$$\mu = \begin{bmatrix} 230 \\ 200 \end{bmatrix},$$

and

$$\Omega = \begin{bmatrix} 10000 & s3000 \\ s3000 & 1400 \end{bmatrix}.$$

The symbol s in front of an entry in the matrix Ω denotes the sign. That is, s is either $+$ or $-$.

Note that the Hessian of the payoff function equals $-b_n\Omega$. Asparouhova, Bossaerts and Plott (2003) noticed (among other things) that the direction of the correlation between a security's excess demand and another security's price was opposite to the sign of the corresponding (off-diagonal) element in the Hessian, or, equivalently, the same as the sign of the corresponding (off-diagonal) element of Ω. To generate strong confirmation of this relation, in the experiments reported on here, the sign s is changed halfway. The aim is to replicate the earlier results (of Asparouhova, Bossaerts and Plott (2003)) in a situation were there is no risk but where quadratic payoff functions replace the (posited) preferences of subjects under uncertainty. If the results are replicable, we expect a switch in the sign of the correlation between a security's excess demand and the other security's price change when we change the sign of the off-diagonal element of the Hessian of the payoff function.

Subjects are assigned one of three levels for the parameter b_n, chosen in such a way as to generate similar pricing (risk premia) as in the uncertainty experiment of Asparouhova, Bossaerts and Plott (2003). See Table 1 for details. Each type also receives a different initial allocation of A and B (nobody receives any Notes to start with, i.e., Notes were in zero net supply). Subjects are not informed of each others' payment schedules or initial holdings, and whether these varied over the course of the experiment (they did not). This way, total endowments of A, B and Notes are unknown, and subjects with knowledge of general equilibrium theory could not possibly compute equilibrium prices and allocations. Therefore, emergence of equilibrium in the experiment cannot be attributed to deliberate application of theoretical concepts.

All accounting is done in terms of an artifical currency, the franc. At the end of the experiment, cumulative earnings are converted to dollar at a pre-announced exchange rate. On average, subjects make about $45 for a three-hour experiment; the range of payments is $0 to approximately $150. These payments, however, inaccurately reflect the size of the incentives during trading. Explicit computations of the amounts of money subjects leave on the table because they did not fully optimize (and assuming that they can trade at closing prices) reveal values over $100 per subject/period in the beginning of an experiment. Some subjects are savvy enough to realize part of these potential gains, but most don't (which explains the substantial range of payouts across subjects). As more subjects realize that there is money to be made, their actions and the ensuing price changes cause these amounts invariably to drop to approximately \$2 in later periods. Subjects seem not to spend the extra effort needed to extract the last couple of dollars.[7]

2.2 Equilibrium Predictions

The competitive equilibrium of the above setup is known as the CAPM. This model makes very specific predictions about prices and allocations. As far as allocations is concerned, CAPM predicts that all subjects will hold securities A and B in the same proportion, a property that is known as *portfolio separation*. In particular, the optimal holding $x_n^*(p)$ of securities A and B by subject n given a vector p of prices of securities A and B (the Notes are taken as numeraire) equals:

$$x_n^*(p) = \frac{1}{b_n}\Omega^{-1}(\mu - p) \qquad (2)$$

Note that this 2×1 vector is proportional to $\Omega^{-1}(\mu - p)$, which is a vector that is common to all subjects. Whence the portfolio separation property. In equilibrium, the common portfolio held by all subjects must of course be the market portfolio (consisting of the entire supply of securities A and B).

The latter immediately generates CAPM's pricing prediction: for all subjects to be willing to hold the market portfolio, it must be priced such that it is optimal, i.e., *the market portfolio is mean-variance efficient*: it provides the best trade-off between the two components of the criterion function in (1), $x \cdot \mu$ (which would be mean payoff if there were uncertainty) and $x \cdot \Omega x$ (the variance of the payoff if there were uncertainty). The explicit pricing formula is:

$$p^* = \mu - B^N \Omega^{-1} \bar{x}^N, \qquad (3)$$

where B^N denotes the harmonic mean of the b_ns across the N subjects,

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

and $\bar{x}^N$ denotes the per-capita supply of risky securities, i.e., the market portfolio:

$$\bar{x}^N = \frac{1}{N} \sum_{n=1}^{N} x_n,$$

(x_n denotes subject n's endowment).

[7] Since each hour consists of approximately three periods, this amounts to \$6 per hour, indicating that subjects need to be given incentives of at least as much in order to generate some evidence of optimizing behavior. The minimum incentive is likely to depend on the nature of the task, but little information on incentive sensitivity in other contexts seems to be publicably available. [Holt and Laury (2002) study the magnitude of risk needed to obtain evidence of risk aversion, however.] Detailed incentive calculations for the experiments at hand can be obtained from the author.

2.3 Disequilibrium Dynamics: The Empirical Evidence

Figure 1 displays the evolution of prices of the three securities (A, B, Notes) in a typical experiment, this particular one ran on 28 May 02. Each observation corresponds to a trade in one of the securities (the price of the remaining securities is taken to be their previous transaction price). Vertical lines delineate periods. Horizontal lines depict equilibrium prices: the top line is the equilibrium price for A, the middle one that for B and the bottom one that of the Notes. There is a shift in equilibrium prices in the fifth period because of the change in the sign of the off-diagonal elements of the Hessian matrix of the payoff function (in experiment 28 May 02, from positive to negative).

While this figure may convey a general impression that transaction prices come close to theoretical predictions and change in the right direction when equilibrium shifts, a few features about equilibration are striking.

- Deviations from equilibrium prices can be as large as 10%. Such deviations are not exceptions and certainly not limited to the first few trades at the beginning of a period and after an equilibrium shift.
- Because of the presence of cash, arbitrage opportunities emerge whenever the price of the Notes deviates from 100. This occurs often in the beginning of a period, when Notes sell for less than 100.
- Prices approach equilibrium from below.
- Prices appear to settle below equilibrium.

These observations are not specific to experiment 28 May 02: they apply to the other experiments as well.

The second observation can readily be explained as the result of a cash-in-advance constraint implicitly imposed in the experiments: all trade takes place against cash, so that a subject who wishes to aggresively buy a security first has to procure enough cash. One easy way to get cash is to short-sell the Notes. Subjects actively engaged in such short-selling of Notes in order to acquire cash to trade the risky securities.

The first observation demonstrates that disequilibrium phenomena must not be ignored. If deviations from equilibrium prices normally amount to about 5% of the price, they are likely to impact inference about equilibrium pricing models from disequilibrium data, even when returns are observed over intervals as long as one month. The typical one-month stock index return, for instance, is about 1% on average, and its volatility 4%. Of course, that in itself does not imply that inference will be biased.

To interpret the third and fourth observations, it is convenient to again view our experimental setting as the certainty equivalent of the CAPM under risk. If risk had replaced nonlinear payoff functions, and subjects exhibited mean-variance preferences in the face of this risk, then the two observations would translate as follows: (•) early on, prices reflect more risk aversion than equilibrium prices, but over time, the risk premium (equity premium) decreases; (•) the equity premium settles above equilibrium levels. The theoretical model we discuss in the next section will provide an explanation for these phenomena. To motivate the theoretical model, we first need to explore the data some more.

Figure 2 depicts the evolution of the distance of the market portfolio from mean-variance efficiency, as measured by the difference in the Sharpe ratio of the market portfolio and the maximal Sharpe ratio (at transaction prices). In our experiments, CAPM describes the competitive equilibrium. At CAPM equilibrium, the market portfolio is mean-variance efficient. But mean-variance efficiency is only a necessary condition for equilibrium. There are other, non-equilibrium price levels for which the market portfolio is mean-variance efficient. Despite evidence that market prices remain as much as 5% below equilibrium (see Figure 1), we see the following.

- The market portfolio quickly moves to the mean-variance frontier and remains there for the remainder of a period.

Figure 3 provides a snapshot of Figure 2: it reveals that the market portfolio is sometimes right on the frontier. We observed the same phenomenon in all the experiments.

Regarding subjects' choices, Figure 4 plots the end-of-period holdings of security A as a proportion of wealth allocated to securities A and B. Each plus sign (+) corresponds to one or more subjects. The weight of A in the market portfolio is indicated with a circle (o). CAPM predicts that all subjects invest in the same portfolio of risky securities; in equilibrium, this common portfolio is the market portfolio. The data do not confirm this. Instead, the following observations can be made.

- Few subjects hold the market portfolio; in fact, portfolio separation fails altogether: very few subjects invest in the same portfolio of securities A and B.

- Subjects' end-of-period holdings are closer to satisfying portfolio separation in general (all hold the same portfolio) and CAPM in particular (all invest in the market portfolio) in the second half of the experiment, when the off-diagonal terms of Ω are negative.

The violation of portfolio separation makes it all the more paradoxical that the market portfolio quickly becomes mean-variance efficient. According to CAPM, optimality of the market portfolio obtains because all subjects want to invest in the same portfolio. Evidently, they do not. Violations of portfolio separation *per se* may not be a puzzle: explicit calculations reveal that subjects leave very little money on the table from not fully diversifying, i.e., from not buying the common optimal portfolio of A and B. After all, diversification has only a secondary effect on subjects' payoffs.[8] It does beg the question, however, why CAPM pricing obtains, because pricing depends crucially on portfolio separation.

[8] Subjects invariably leave less than one dollar per subject/period on the table from not diversifying. Explicit calculations can be obtained from the author. Subjects would often lose money if they were to invest in the common optimal portfolio. This occurs when they are not holding the optimal amount of Notes given their level of b_n.

The challenge to our theory of dynamics will be to explain this paradox. It is not specific to the experiment 28 May 2002, having been recorded in the other experiments as well (and in all uncertainty experiments – see Bossaerts, Plott and Zame (2002)).

The last feature of our experiments directly concerns the origin of price changes. Using data from numerous asset market experiments, Asparouhova, Bossaerts and Plott (2003) documents that transaction price are correlated not only with a security's own theoretical excess demand, but also with other securities' theoretical excess demand. Table 2 demonstrates that this phenomenon is replicable; it is recorded in each of our certainty-equivalent experiments as well.

- A security's transaction price changes are significantly correlated with excess demand in other securities.

Note that the correlation emerges despite the lack of physical linkages between markets. That is, there is nothing in the microstructure of our markets that would spuriously induce correlation: orders cannot be made conditional on events in other markets, payments cannot be made in other securities, etc.

Table 2 reveals a remarkable link between the slopes in projections of transaction prices onto theoretical excess demands, on the one hand, and the corresponding element of Ω, on the other hand, thus confirming another aspect of the findings in Asparouhova, Bossaerts and Plott (2003).

- The correlation between a security's transaction price changes and another security's excess demand is related to the corresponding off-diagonal element in the Hessian of the payoff function.

In periods 1 to 4 of experiment 25 Apr 02, for instance, Ω equals:

$$\Omega = \begin{bmatrix} 10000 & 3000 \\ 3000 & 1400 \end{bmatrix}.$$

Compare this to the matrix of slope coefficients in the projections of the transaction prices of securities A and B onto the theoretical excess demands of A and B:

$$\begin{bmatrix} 0.261 & 0.080 \\ 0.078 & 0.043 \end{bmatrix}.$$

These slope coefficients are significant and their sign and relative magnitude are comparable to the sign and relative magnitude of the corresponding element of Ω. In periods 5 to 8 of the same experiment, Ω equals:

$$\Omega = \begin{bmatrix} 10000 & -3000 \\ -3000 & 1400 \end{bmatrix}.$$

The slope coefficients of the projections of transaction price changes onto theoretical excess demands generate the following matrix:

$$\begin{bmatrix} 0.308 & -0.095 \\ -0.433 & 0.162 \end{bmatrix}.$$

Notice how the off-diagonal elements of both matrices flipped signs compared to the first four periods.

Before we embark on modeling, it should be emphasized that the above phenomena are based on noisy observations, creating the possibility that they are masking more fundamental phenomena. One of the roles of theory is to determine whether we have been looking at the rights things, e.g., whether the above correlations aren't spurious because the variables at hand are proxying for other, more fundamental ones. We will find a number of instances where this is indeed the case.

3 A Model Of Equilibration

We now develop a model of equilibration. The goal is to predict price and allocation dynamics in a competitive, continuous trading system, where all transactions are small and determined without the intervention of some central price or quantity setting institution. It only relies on one "axiom:" that order imbalance leads to price pressure; upward if there are more offers to buy (bids) than offers to sell (asks); negative otherwise.

The behavioral foundation of the model is unlike in standard economic theory. Agents are not assumed to fully optimize. Instead, they adjust their portfolio holdings gradually, into the direction of maximum increase in utility. Prices move as a response to orders that reflect this local optimization. Prices will therefore not change (directly) as a result of excess demand. There is an indirect relationship, however, which is simple for the case of the CAPM, and happens to be the one found in the data.

Assume a pure exchange economy. Let there be N (types of) agents, indexed $n = 1, ..., N$. There are $I + 1$ securities, indexed $i = 1, ..., I + 1$. Security $I + 1$ will be the numeraire - think of it as cash and/or Notes. Agent n holds a (vector) x_n of the first I security and h_n of security $I + 1$. Utility U_n is separable in choices of the first I securities x and choice h of security $I + 1$, strictly concave in x and linear in h. (This can readily be relaxed; here, we want to get to the CAPM in particular.) Write:

$$U_n(x, h) = u_n(x) + h.$$

Agents maximize utility. They take prices as given.

In standard competitive analysis, agents are assumed to choose quantities x_n^* and h_n^* that maximize their utility, subject to a budget constraint evaluated at equilibrium prices p^*, namely

$$p^* \cdot x_n^* + h_n^* = p^* \cdot x_n + h_n.$$

p^* is determined such that total demand equals total supply. Note that x_n^* is set such that

$$p^* = \nabla u_n(x_n^*). \tag{4}$$

We'll need this later.

We do not follow the route of standard competitive analysis. Instead, we start from the following behavioral postulate. As before, let x_n and h_n denote current holdings of agent n. (S)he submits orders z_n. These are infinitesimal and *change her/his holdings in the direction of maximum gain in utility.* Orders have to be feasible, though: they have to be within the agent's budget – whether they are is determined by the prices that the agent sends along with the orders.[9] We assume that all agents send in the same set of prices p. p represent the state of the market.[10] We'll specify below how p is determined. The fact that all agents submit orders at the same p reflects competition. It represents the price of executable orders, which individual agents cannot influence.

Utility trade-offs along agent n's budget constraint can be described as a function L_n of x, with

$$L_n(x) = U_n(x, p \cdot (x - x_n) + h_n).$$

The orders that maximizes the local gain in utility are elements of the gradient of L_n:

$$z_n = \nabla L_n(x_n).$$

Effectively, we are assuming that agents submit an order in the direction that maximizes the directional derivative of utility constrained by the budget, subject to a quadratic adjustment cost. That is, z_n solves:

$$\max_z \left\{ \nabla L_n(x_n) \cdot z - \frac{1}{2} z \cdot z \right\}.$$

Quasi-linearity of the utility function implies:

$$z_n = \nabla u_n(x_n) - p. \tag{5}$$

The orders z_n are sent to the marketplace, where they add up to a vector of order imbalances Z:

$$Z = \sum_{n=1}^{N} z_n.$$

In general, Z is different from zero, so not all orders can be executed. A number of them (sometimes none) will be executed, on the basis of some screening device which we leave unspecified for the moment. The order imbalance creates proportional price pressure. That is, the subsequent change in price, Δp, equals:

$$\Delta p \sim Z = \left[\sum_{n=1}^{N} \nabla u_n(x_n) - p \right]. \tag{6}$$

[9] Note that the experimental software does check for budget feasibility whenever an order is submitted.

[10] Agents may submit orders at different prices as well, but these won't be executable (if extra-marginal), in which case we ignore them, or they will be executed at p anyway (if infra-marginal), in which case we imagine that the agent submitted the price p.

This generates a set of difference equations which describe the evolution of prices, conditional on holdings x_n. If, in addition, we specify how x_n changes, i.e., if we specify which orders get executed, we have a more complex system of difference equations that describes the simultaneous movement of prices and allocations.

We now explore the implications of the model, first for pricing, and next for allocations, and relate these to the experimental findings.

4 Implications for Pricing

4.1 Relation between Price Changes and Excess Demands

We first demonstrate that price changes will be related to excess demands with a projection matrix that is proportional to Ω, as in the data.

First consider the following linear approximation of the gradient function around the globally optimal holdings x_n^* (given prices p):

$$\nabla u_n(x_n) \approx \nabla u_n(x_n^*) - Hu_n(x_n^*)(x_n^* - x_n),$$

where H is the hessian of u_n. Also, remember that x^* is determined by the following set of equations:

$$p = \nabla u_n(x_n^*).$$

Use these to re-arrange the expression for the orders z_n [see (5)]:

$$z_n = \nabla u_n(x_n) - \nabla u_n(x_n^*) = -Hu_n(x^*)(x_n^* - x_n).$$

Therefore,

$$\Delta p \sim Z = -\sum_{n=1}^{N} Hu_n(x_n^*)(x_n^* - x_n).$$

In the case of the CAPM,

$$Hu_n(x_n^*) = -b_n\Omega$$

[see (1)] and the above approximation of the gradient function is without error. Thus,

$$\Delta p \sim \Omega \sum_{n=1}^{N} b_n(x_n^* - x_n) \tag{7}$$

The matrix Ω provides the link between price changes and excess demands.

Note however, that price changes are actually not related to total excess demand, but *to risk-aversion-weighted individual excess demands.* In the empirical analysis of equilibration dynamcis, we projected price changes onto the former, not the latter. But obviously the two variables are very highly correlated. In other words, when we discover that the matrix of slope coefficients in projections of price changes onto total excess demand reveals the same structure as the Hessian matrix, this emerges because total excess demand correlates highly with a more fundamental variable, namely, risk-aversion-weighted individual excess demand. When re-running the projections, using risk-aversion-weighted individual excess demands as explanatory variables, only marginal improvements were obtained, as expected.

4.2 Mean-Variance Optimality of the Market Portfolio

To study how far the market portfolio remains below the mean-variance efficient frontier during equilibration, we simplify things, assuming that prices settle before trade takes place, and we focus on the optimality of the market portfolio at transaction prices only. That is, we study efficiency of the market portfolio at prices for which local excess demands equilibrate. While it is doubtful that this describes trade in the experiments, it may provide an adequate model of field markets, which are populated with professional traders who should be able to ensure that their orders are executed at (local) equilibrium only. Consequently, the analysis to follow may be more relevant to field markets than to experimental markets. Nevertheless, it generates predictions that are remarkably in agreement with the experimental data.

Mathematically, we will be exploring the steady-state solution(s) of the dynamics in (6) for a given set of initial holdings $\{x_1, ..., x_n, ..., x_N\}$, subsequently altering the holdings at the steady state prices (at which point local excess demands equilibrate), computing new steady state prices, to be used to change the holdings again, etc. The dynamics of prices and allocations will thus be described by a complex set of interconnected difference equations.

Local equilibrium prices are steady-state solutions to (6), which we repeat here for convenience:

$$\Delta p = \lambda \left[\sum_{n=1}^{N} \nabla u_n(x_n) - p \right],$$

where λ is some (positive) constant of proportionality. In the case of CAPM preferences, this system of equations boils down to:

$$\Delta p = \lambda \left[-\Omega(\sum_{n=1}^{N} b_n x_n) + \mu - p \right]. \tag{8}$$

Its (single) steady-state solution equals:

$$p^o = \mu - \Omega \frac{1}{N} \sum_{n=1}^{N} b_n x_n. \tag{9}$$

At these prices, it is easy to identify the mean-variance optimal portfolio of risky securities, namely:

$$\frac{1}{N} \sum_{n=1}^{N} b_n x_n. \tag{10}$$

We refer to this portfolio as the *risk-aversion-weighted endowment (RAWE) portfolio.*

Equation (9) essentially states that orders will be in balance when the RAWE portfolio is mean-variance optimal. All agents will then be able to implement the adjustments they set out to make to their portfolios. At the prices in (9), the adjustments provide maximum local gain in utility, subject to a quadratic adjustment cost. Once the adjustments are made, agents may wish to return to the market to make further adjustments, but these will generally be at different prices.

The RAWE portfolio is closely related to the market portfolio. When risk aversion is independent of holdings, the two portfolios are proportional, which implies that they put the same weight on each asset (in a portfolio-theoretic sense, they are identical):

$$\frac{1}{N}\sum_{n=1}^{N} b_n x_n = (\frac{1}{N}\sum_{n=1}^{N} b_n)\frac{1}{N}\sum_{n=1}^{N} x_n = (\frac{1}{N}\sum_{n=1}^{N} b_n)\bar{x}^N.$$

At CAPM equilibrium the two also coincide, because holdings are proportional to the market portfolio:

$$\frac{1}{N}\sum_{n=1}^{N} b_n x_n = \frac{1}{N}\sum_{n=1}^{N} b_n \left[\frac{B^N}{b_n}\left(\frac{1}{N}\sum_{n=1}^{N} x_n\right)\right] = B^N \bar{x}^N.$$

As a consequence, if the equilibration process starts out at endowments that are independent of risk aversion, *prices immediately move to levels for which the market portfolio is mean-variance optimal.* As holdings change, the market portfolio may become sub-optimal, but it eventually will have to move back to the mean-variance frontier.[11]

The weights in the market portfolio and the RAWE portfolio can be expected to be close, which means that the market portfolio remains close to the mean-variance efficient frontier since RAWE stays on it during the entire equilibration process. In the experiments, the two are close. Confirmation of our model comes from the finding that the market portfolio indeed quickly moves to the frontier at the beginning of a period, and remains close to it from then on, as reported in the empirical section.

In one uninteresting case, the market portfolio and RAWE always coincide, namely, when agents have the same risk aversion coefficient b_n. It would imply that the market portfolio remains optimal throughout equilibration. This is unlikely to provide an accurate description of field markets, and indeed, the market portfolio (or reasonable proxies), while historically close to optimal, can nevertheless be improved upon – see, e.g., Fama and French (1992). In the certainty-equivalent experiments we discussed in the empirical section, we did vary the paramater b_n across subjects. See Table 1.

[11] The arguments in Smale (1976) imply that only Pareto-optimal allocations are limit points of our equilibration process. Because CAPM relies on constant absolute risk aversion, there are no wealth effects, and hence, all Pareto-optimal allocations are supported by the same prices, namely, CAPM equilibrium prices. Hence, limit points to our price dynamics must be CAPM equilibrium prices, at which the market portfolio is mean-variance optimal.

4.3 The Equity Premium

While the market portfolio will be close to mean-variance efficient by virtue of its association with the RAWE portfolio, it may be priced very differently from CAPM equilibrium predictions. To be more precise, the equity premium (the expected return on the market portfolio minus the risk free rate) may be very different from the equilibrium equity premium. To put it in yet other terms, the risk aversion implicit in securities prices may be very different from that reflected in CAPM equilibrium prices.

To understand this, remember that, at CAPM equilibrium, securities prices p^* reflect the *harmonic mean risk aversion* B^N of the agents:

$$p^* = \mu - B^N \Omega \bar{x}^N$$

[see (3)]. In our equilibration dynamics, and provided endowments initially are independent of risk aversion (which implies that the market portfolio will be mean-variance optimal initially), prices reflect the *arithmetic mean risk aversion* instead:

$$p^o = \mu - \Omega \frac{1}{N} \sum_{n=1}^{N} b_n x_n = \mu - \left(\frac{1}{N} \sum_{n=1}^{N} b_n \right) \Omega \frac{1}{N} \sum_{n=1}^{N} x_n = \mu - \left(\frac{1}{N} \sum_{n=1}^{N} b_n \right) \Omega \bar{x}^N.$$

Because the arithmetic mean is always strictly higher than the harmonic mean (except in the uninteresting case where b_n is the same across n), the risk aversion implicit in securities prices early on in our price process will be higher than at equilibrium. Put differently, the *equity premium will initially be higher than when equilibrium is reached.* The experimental data confirm this prediction: prices approach equilibrium levels from below.

5 Implications For Allocations

In our equilibration model, prices move faster than trades, so that all trades take place at local equilibrium prices. If we are willing to assume that a fixed fraction of local demands are filled every time there is trade, we generate a simple theory of allocation dynamics.

Upon trading, agent n changes his holdings from x_n to x'_n, where

$$x'_n = x_n + \delta \nabla L_n(x_n).$$

δ denotes the fraction of (local) demand that is filled each period. In the case of CAPM, this boils down to the following:

$$x'_n = x_n - \delta \Omega \left[b_n x_n - \frac{1}{N} \sum_{\nu=1}^{N} b_\nu x_\nu \right].$$

Let $\Delta x_n = x'_n - x_n$. Re-writing the above equation, we obtain:

$$\Delta x_n = -\delta\Omega \left[b_n x_n - \frac{1}{N} \sum_{\nu=1}^{N} b_\nu x_\nu \right]. \tag{11}$$

In words: trade (Δx_n) is a linear transformation ($\delta\Omega$) of the deviation between subject n's risk-aversion-weighted portfolio ($b_n x_n$) and the RAWE portfolio ($\frac{1}{N} \sum_{\nu=1}^{N} b_\nu x_\nu$).

It is interesting to study what such dynamics imply in the extreme case where all subjects start out with the same portfolio (the market portfolio). In that case,

$$\left[b_n x_n - \frac{1}{N} \sum_{\nu=1}^{N} b_\nu x_\nu \right] = \left(b_n - \frac{1}{N} \sum_{\nu=1}^{N} b_\nu \right) \frac{1}{N} \sum_{\nu=1}^{N} x_\nu = \left(b_n - \frac{1}{N} \sum_{\nu=1}^{N} b_\nu \right) \bar{x}^N.$$

Hence, agents trade a portfolio that is a linear transformation of a portfolio proportional to the market portfolio. Except in the unlikely event that the market portfolio is an eigenvector of the matrix Ω defining the transformation, the new holdings will not be proportional to the market portfolio. That is, agents trade away from the market portfolio, and consequently, end up with holdings that *violate portfolio separation*.

The mathematical equations suggest in what direction portfolios change. Imagine that Ω is diagonal, the diagonal elements of Ω being the payoff variances. In that case, volume (the absolute value of the elements in Δx_n) will be highest for the high-variance securities. That is, most portfolio adjustments take place in the high-variance securities. The sign of the changes in an agent's holdings of risky securities depends on his risk aversion (b_n) relative to the average risk aversion ($\frac{1}{n} \sum_\nu b_\nu$). More risk averse agents sell risky securities (the entries of Δx_n are negative); less risk averse agents buy. Effectively, more risk averse agents unload risky securities, paying more attention to the most risky securities, because that way their local gain in utility is maximized. Likewise, less risk averse agents do what is locally optimal: increase risk exposure by buying the most risky securities first.

When Ω is non-diagonal, the sign of the off-diagonal elements interferes with the above dynamics. Intuitively, when the off-diagonal elements are negative, i.e., when the payoff covariances are negative, securities are natural hedges for each other, and the market portfolio provides diversification. Increasing one's risk exposure by buying mostly risky securities (or decreasing one's risk exposure by selling mostly risky securities) leads to a less diversified portfolio, i.e., to utility losses. Maximum local gains in utility are instead obtained by trading combinations of securities that are closer to the market portfolio. As a consequence, agents' portfolios of risky securities will remain closer to the market portfolio than if payoff covariances were zero (or positive, for that matter). That is, we expect to see less extreme violations of portfolio separation.

Eventually, equilibrium will be reached, at which point portfolio separation would be restored. But agents must be willing to implement all trades necessary to reach equilibrium. At a certain point in the equilibration process, agents may not perceive enough gains to cover the effort it takes to trade. Trading halts before portfolio separation is restored.

The intuition behind this result is simple. It is well known that the gains from diversification are second-order. Agents first of all should ensure that they hold the right mix of risky and riskfree securities. Once they acccomplish this, further rebalancing of their portfolio to fine-tune diversification will lead to utility gains, but they are an order of magnitude smaller, and hence, ignored if too small.

We observe this in the experiments. End-of-period holdings indeed violate portfolio separation (see Figure 4). But our theory makes finer predictions: violations of portfolio separation should be less pronounced if the off-diagonal element of Ω is negative. This is upheld in the data. Subjects' holdings of risky securities are closer to portfolio separation in the first half of experiment 28 May 02. The same conclusion obtains in the other experiments. In experiment 28 Nov 01, however, the off-diagonal element of Ω is positive in the first half. There, the finding that portfolio separation violations are more pronounced in the first half could have been attributed to learning (subjects take time to appreciate the benefits of diversification). This possibility is rejected, of course, because of the findings in other experiments, where portfolio separation violations are bigger in the second half.

It is interesting to explore the effects of allocation dynamics on pricing. Assume covariances (off-diagonal elements of Ω) are zero (or positive). More risk averse agents trade towards portfolios that put less weight on securities with high payoff variance than the market portfolio; *vice versa*, more risk tolerant agents trade towards portfolios that overweigh the most risky securities. Consequently, the RAWE portfolio (i.e., the risk-aversion-weighted endowment portfolio) will be overweighted in low-risk securities and underweighted in high-risk securities, relative to the market portfolio. Since the RAWE portfolio is mean-variance optimal throughout equilibration, this means that the return on the market portfolio can be improved upon by combining it with a zero-investment portfolio long in low-variance securities and short in high-variance securities. We will come back to this point shortly, because it may provide an original interpretation of historical experience in field markets.

6 Implications for Tests of CAPM on Field Data

The goal of this paper is to investigate the implications of our model for tests of the CAPM on field data. We assume, perhaps rightly so, that field markets equilibrate faster than our experimental markets, so that we can ignore transactions that occur at prices away from (local) equilibrium. The presence of professional traders in field markets, absent in our experimental markets, may motivate this assumption.

Our theoretical analysis predicts the following behavior of prices and allocations. First, transaction prices will be such that the RAWE portfolio is continuously on the frontier. The market portfolio may be close to the frontier but only when it is a good proxy for RAWE. Second, risk aversion in transaction prices may seem higher than actual risk aversion among investors, or, equivalently, the equity premium may be larger than expected given investors' risk aversion. Third, even if all investors at times hold the market portfolio, they have the tendency to trade away from it once parameters change (endowments, expected payoffs, variances and covariances, etc.) and may not have enough incentives to trade all the way back to the market portfolio. Additional conclusions can be drawn about the composition of RAWE relative to the market portfolio during the equilibration process, provided we know the signs of the payoff covariances.

It is best to illustrate these findings with a simulation of a market where agents all start with the same portfolio and trade on the basis of our principles. To make cross-reference with the experiments easy, all other parameters are taken to be the same as in the first half of experiment 28 May 02 (μ, Ω, distribution of b_n among agents, etc.). δ, the fraction of total (local) demand that is filled each period, equals 1%.

Figure 5 plots the predicted evolution of the transaction prices. Initially, prices deviate substantially (up to 40%!) from equilibrium prices (indicated with horizontal lines). Equilibrium is reached from below, i.e., risk implicit in transaction prices is higher than at CAPM equilibrium. The top panel of Figure 6 confirms that the equity premium is initially much higher than at equilibrium; it decreases only gradually over time. Figure 6 also displays the evolution of the distance from mean-variance optimality of both the RAWE and the market portfolio. This distance is measured as the difference between the Sharpe ratio of the portfolio (RAWE; market) and the maximum possible Sharpe ratio. As predicted, the RAWE portfolio is constantly on the mean-variance frontier. The market portfolio starts on the frontier, and then moves below it. But it never moves far below the frontier, because it tracks the RAWE portfolio. It eventually returns to the frontier - evidently very slowly, though.

Figure 7 displays the evolution of the mean trade size (number of units of each security per agent). After about 60 periods, agents exchange only tiny quantities compared to early trading, indicating that most gains from trade are already exhausted. If agents halt trading after 60 periods because they perceive little gains, they end up with portfolios that grossly violate portfolio separation. Figure 8 confirms this. It plots the evolution of the cross-sectional standard deviation of the proportion of wealth in securities A and B that agents allocate to security A. Portfolio separation obtains when this standard deviation equals zero. Such is the case at the start (they all hold the same portfolio). Agents quickly move away from the market portfolio. The standard deviation rises to close to 0.5, and decreases only gradually afterwards. After 60 periods, it is still above 0.4.

All this implies the following for tests of the CAPM in markets that may not be in equilibrium when price and holding data are collected. Foremost, the analysis indicates that tests of the core CAPM pricing prediction (that the market portfolio is mean-variance efficient) still make sense. More specifically, the following can be said.

- Provided the market portfolio is a good proxy for the RAWE portfolio, tests of CAPM's core pricing prediction allows one to infer whether CAPM provides a correct model of the pricing of risk in the economy.
- Such a test makes sense even if CAPM's portfolio holding predictions (portfolio separation) are grossly violated.
- One expects the equity premium to over-estimate agents' risk aversion, however.

But to what extent is the market portfolio a good proxy for RAWE? Evidence from field markets has suggested that the market portfolio is *not* mean-variance efficient, but instead can be improved upon by combination with (at least) one zero-investment portfolio. This raises the issue of whether this empirically motivated, adjusted market portfolio is a better proxy for RAWE (which, in our theory at least, is mean-variance optimal throughout equilibration) than the market portfolio itself.

The RAWE portfolio can obviously trivially be decomposed into the market portfolio and a zero-investment adjustment portfolio. The latter accommodates overweighing and underweighing of investments in RAWE relative to the market portfolio. Mean-variance efficiency of the RAWE portfolio then translates into mean-variance efficiency of a combination of the market portfolio and this adjustment portfolio. Since mean-variance efficiency implies a linear relationship between expected returns and "betas" of the portfolios that constitute the mean-variance efficient frontier, our analysis implies a multi-factor asset pricing model for the cross-section of mean returns. In it, the market and the zero-investment adjustment portfolio are the "factor portfolios."

The allocation dynamics in our equilibration model determine the composition of the adjustment portfolio. If the covariance between the payoffs of all securities is non-negative, and if agents are initially endowed with the market portfolio, then the adjustment portfolio accommodates the ensuing overweighing in RAWE of securities with low payoff variance and underweighing of securities with high payoff variance. That is, the adjustment portfolio will be long securities with low payoff variance, and short securities with high payoff variance. This may provide an explanation for why investment performance could have been improved historically by overweighing low-volatility stock and underweighing high-volatility stock. See Haugen and Baker (1996).

If common stock in high-value firms can be labeled as low-variance securities, and growth stock as high-variance securities, we essentiall have obtained an intriguing explanation of the "value effect" in historical (field market) returns as well.[12] As Fama and French (1992) and others have shown, a combination of (a proxy of) the market portfolio and a zero-investment portfolio long in value stock and short in growth stock outperforms the market portfolio in mean-variance terms. The composition of this superior portfolio is thus the same as would be our RAWE during equilibration: overweight in the low-variance securities and underweight in the high-variance securities. And RAWE outperforms the market, because it is always mean-variance optimal (at least in theory).[13,14]

[12]The return volatility on high (book-to-market) value stocks is lower than that of low value (i.e., growth) stocks, although the discrepancy is more pronounced among small firms. See, e.g., Fama and French (1995).

[13]Fohlin (2000) documents that Germany experienced a reverse value effect in the period 1881-1913, which means that, to improve upon the market portfolio, one should have underweighted value (high book-to-market) stock and over-weighted growth stock. Interestingly, high book-to-market stock in that period was predominantly stock in small firms, which were exceptionally volatile. See their Table II. Such a finding is consistent with our model: RAWE would put less weight on the high book-to-market securities, because they are more volatile.

[14]Field market research has also identified a size effect: the optimal portfolio puts more weight on small firms than the market portfolio and less weight on large firms. See also Fama and French (1992). The size effect has been documented to often disappear, however,

7 Conclusion

Perhaps the most important contribution of this paper is that, in spite of disequilibrium, one specific portfolio remains continuously on the mean-variance efficient frontier, namely the risk-aversion-weighted endowment (RAWE) portfolio. Only to the extent that RAWE and the market portfolio are linked can tests of mean-variance efficiency of the market portfolio be defended as tests of the CAPM. Absent information on the precise link between the two, however, it would be much cleaner to directly work with the RAWE portfolio. In other words, empiricists may have been looking at the wrong portfolio. Still, recent interest in adjustments to the market portfolio based on book value may be interpreted as the beginning of a quest for the RAWE portfolio, because the adjustments seem to be going in direction suggested by this paper's theory. In any event, the paper does underscore that it makes sense to test the CAPM (or related multi-factor models) even if its predictions about portfolio holdings are grossly violated and even if the equity premium looks excessively high. Both could merely be effects of disequilibrium trading.

References

Anderson, C., S. Granat, C. Plott and K. Shimomura (2003), "Global Instability in Experimental General Equilibrium: The Scarf Example," *Journal of Economic Theory*, forthcoming.

Arrow, K. and F. Hahn (1971), *General Competitive Analysis*, San Francisco: Holden-Day.

Asparouhova, E. and P. Bossaerts (2002), "Price Discovery In Financial Markets: Analysis Of Two Systems Of Differential Equations," Caltech working paper.

Asparouhova, E., P. Bossaerts and C. Plott (2003), "Excess Demand And Equilibration In Multi-Security Financial Markets: The Empirical Evidence," *Journal of Financial Markets* 6, 1–22.

Benink, H. and P. Bossaerts (2001), "An Exploration of Neo-Austrian Theory Applied To Financial Markets," *Journal of Finance* 54, 1011-1028.

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

Bossaerts, P., C. Plott and W. Zame (2003), "The Evolution of Prices and Allocations in Markets: Theory and Experiment," Caltech/UCLA manuscript.

suggesting that the effect may have been the result of hindsight bias. If genuine, however, it would be hard for our theory to explain it. Effectively, one has to argue that shares in small companies (temporarily) end up in the hands of the more risk averse agents.

Fama, E. and K. French (1992), "The Cross–Section of Expected Stock Returns," *Journal of Finance* 47, 427–465.

Fama, E. and K. French (1995), "Size and Book-to-Market Factors in Earnings and Returns," *Journal of Finance* 50, 131–155.

Fohlin, C. (2000), "Has the Cross-Section of Average Returns always been the Same? Evidence from Germany, 1881-1913," Caltech working paper.

Haugen, R. and N. Baker (1996), "Commonality in the Determinants of Expected Stock Returns," *Journal of Financial Economics* 41, 401–439.

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

Plott, C. (2000), "Market Stability: Backward-Bending Supply in a Laboratory Experimental Market," *Economic Inquiry* 38, 1-18.

Smale, S. (1976), "Exchange Processes with Price Adjustment," *Journal of Mathematical Economics* 3, 211–226.

Table 1: Experimental Design Data.

Exp.[a]	Subject Cat. (#)[c]	b_n[b] $\times 10^{-3}$	Signup Reward (franc)	Endowments A	Endowments B	Endowments Notes	Cash (franc)	Loan Repayment (franc)	Exchange Rate $/franc
28 Nov 01	14	2.30	125	2	8	0	400	2340	0.06
	14	0.28	125	8	2	0	400	2480	0.06
	14	0.15	125	2	8	0	400	2365	0.06
20 Mar 02	10	2.30	125	2	8	0	400	2320	0.06
	10	0.28	125	8	2	0	400	2470	0.06
	10	0.15	125	2	8	0	400	2370	0.06
24 Apr 02	14	2.30	125	2	8	0	400	2320	0.06
	13	0.28	125	8	2	0	400	2470	0.06
	13	0.15	125	2	8	0	400	2370	0.06
28 May 02	13	2.30	125	2	8	0	400	2320	0.06
	12	0.28	125	8	2	0	400	2470	0.06
	12	0.15	125	2	8	0	400	2370	0.06

[a]Date of experiment.
[b]Coefficient in the payoff function (1) assigned to subject.
[c]Number per subject type.

Table 2: Projections Of Transaction Price Changes Onto Theoretical Excess Demands

Experiment	Periods	Sign[a]	Security	Coefficients[b] Intercept	Excess Demand A	Excess Demand B	R^2	F-statistic[c]	N[d]	ρ[e]
28 Nov 01	1-4	$-$	A	0.004	0.476	-0.154	0.08	49.5	1142	-0.08^{**}
				(0.051)	(0.049)	(0.016)		(< .01)		
			B	-0.276	-0.040	0.027	0.08	20.7	1142	-0.00
				(0.052)	(0.049)	(0.016)		(< .01)		
	5-8	$+$	A	-0.948	0.445	0.127	0.09	57.7	1225	-0.25^{**}
				(0.167)	(0.044)	(0.014)		(< .01)		
			B	-0.243	0.003	0.013	0.04	28.4	1225	0.03
				(0.102)	(0.027)	(0.001)		(< .01)		
20 Mar 02	1-4	$+$	A	-0.250	0.211	0.058	0.03	11.6	667	0.01
				(0.111)	(0.045)	(0.013)		(< .01)		
			B	-0.717	-0.146	0.003	0.20	80.9	667	0.04
				(0.173)	(0.070)	(0.020)		(< .01)		
	5-8	$-$	A	-0.538	0.562	-0.178	0.10	27.3	487	-0.03
				(0.113)	(0.081)	(0.028)		(< .01)		
			B	-0.246	-0.196	0.090	0.12	33.8	487	0.01
				(0.065)	(0.046)	(0.016)		(< .01)		
25 Apr 02	1-4	$+$	A	-0.674	0.261	0.080	0.05	15.0	612	-0.08^{*}
				(0.194)	(0.045)	(0.017)		(< .01)		
			B	-0.762	0.078	0.043	0.06	20.4	612	0.05
				(0.180)	(0.046)	(0.015)		(< .01)		
	5-8	$-$	A	-0.394	0.308	-0.095	0.05	18.9	676	0.00
				(0.114)	(0.070)	(0.024)		(< .01)		
			B	-0.326	-0.433	0.162	0.13	50.6	676	-0.04
				(0.081)	(0.050)	(0.018)		(< .01)		
28 May 02	1-4	$+$	A	-0.386	0.200	0.050	0.05	20.2	819	0.04
				(0.278)	(0.040)	(0.014)		(< .01)		
			B	-0.999	0.056	0.039	0.10	45.1	819	0.00
				(0.159)	(0.028)	(0.008)		(< .01)		
	5-8	$-$	A	-0.148	0.072	-0.021	0.01	3.9	561	-0.14^{**}
				(0.088)	(0.030)	(0.010)		(0.02)		
			B	-0.354	-0.060	0.035	0.07	22.0	561	0.00
				(0.085)	(0.029)	(0.010)		(< .01)		

[a] Sign s of the off-diagonal element of the matrix Ω defining the Hessian of the payoff function:

$$\Omega = \begin{bmatrix} 10000 & s3000 \\ s3000 & 1400 \end{bmatrix}.$$

The OLS coefficient matrix evidently inherits the structure (including the sign s) of this matrix.

[b] OLS projections of transaction price changes onto (i) an intercept, (ii) the theoretical excess demands for the two risky securities (A and B). Time advances whenever one of the three assets trades. Standard errors in parentheses.

[c] p-level in parentheses.

[d] Number of observations.

[e] Autocorrelation of the error term; * and ** indicate significance at the 5% and 1% level, respectively.

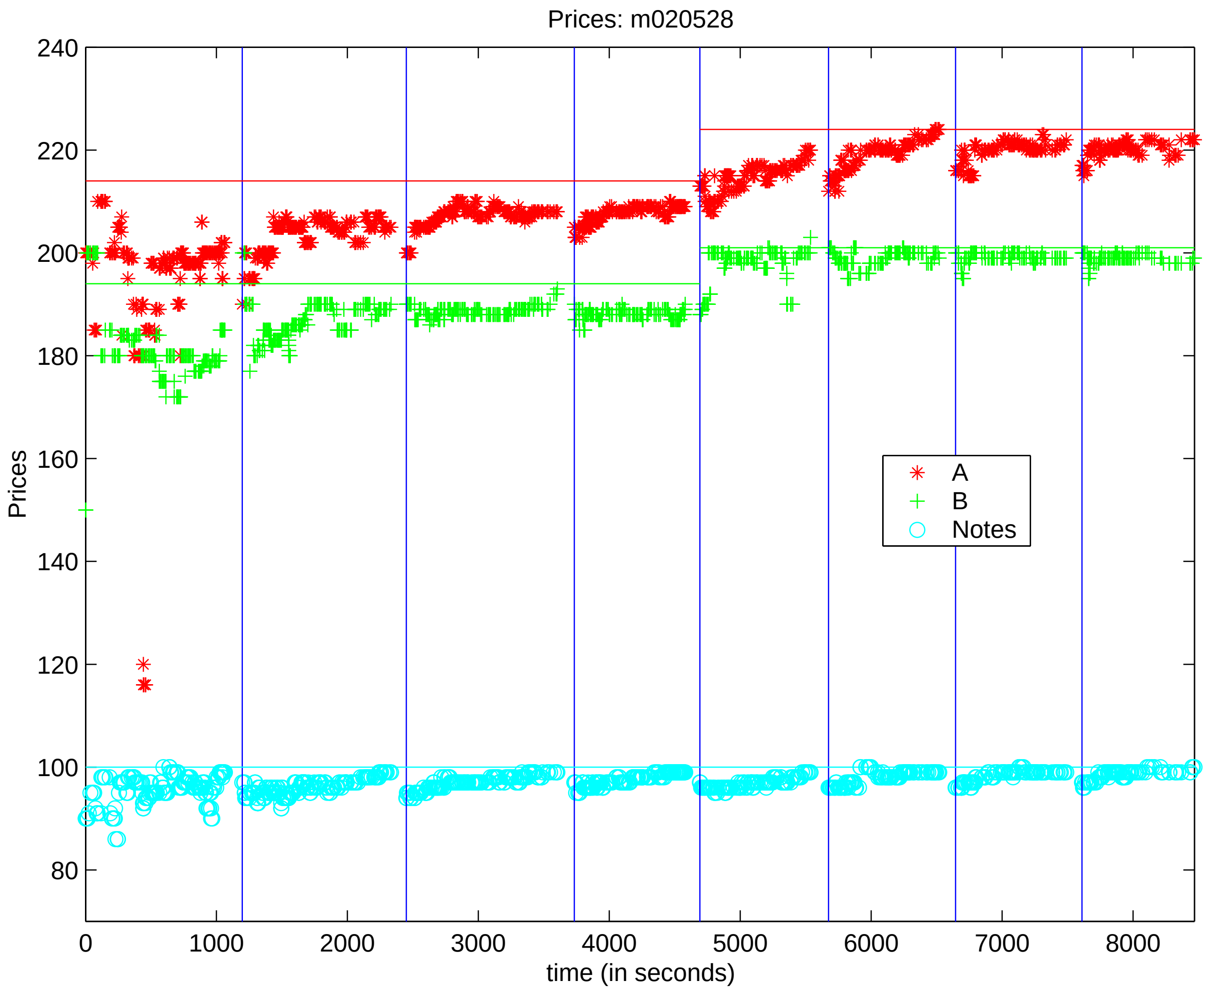

Figure 1: Plot of the evolution of transaction prices in experiment 28 May 02. Vertical lines delineate periods. Horizontal lines depict equilibrium prices. The equilibrium in periods 1-4 is different from that in periods 5-8 because of a change in the sign of the off-diagonal element of the Hessian of the payoff function.

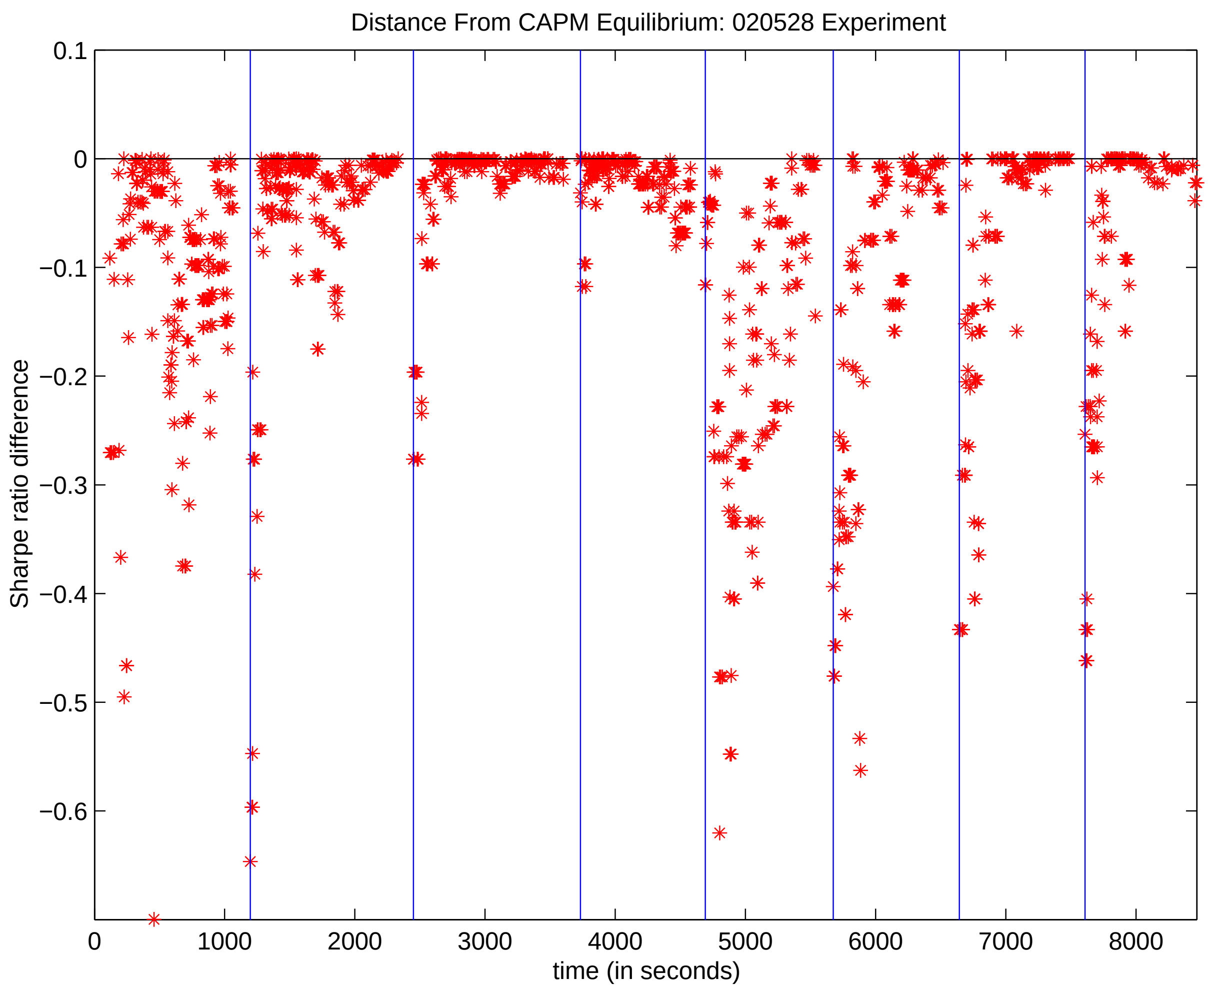

Figure 2: Plot of the evolution of the distance between the market portfolio and themean-variance efficient frontier, as measured by the difference between the Sharpe ratio of the market portfolio and the maximal Sharpe ratio, experiment 28 May 02. Vertical lines delineate periods.

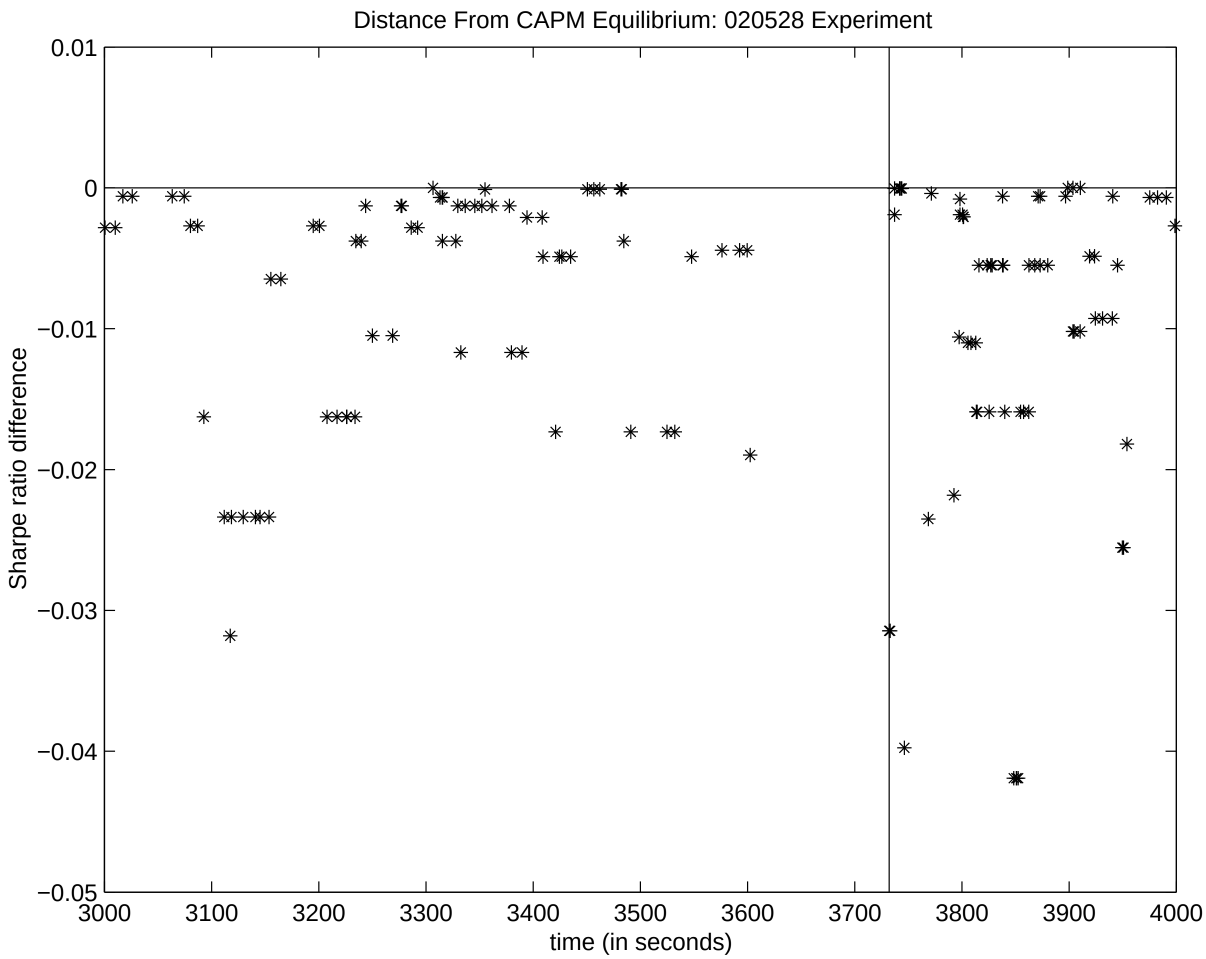

Figure 3: Snapshot of previous figure (plot of the evolution of the distance between the market portfolio and themean-variance efficient frontier, as measured by the difference between the Sharpe ratio of the market portfolio and the maximal Sharpe ratio, experiment 28 May 02). Vertical line indicates the end of period 3.

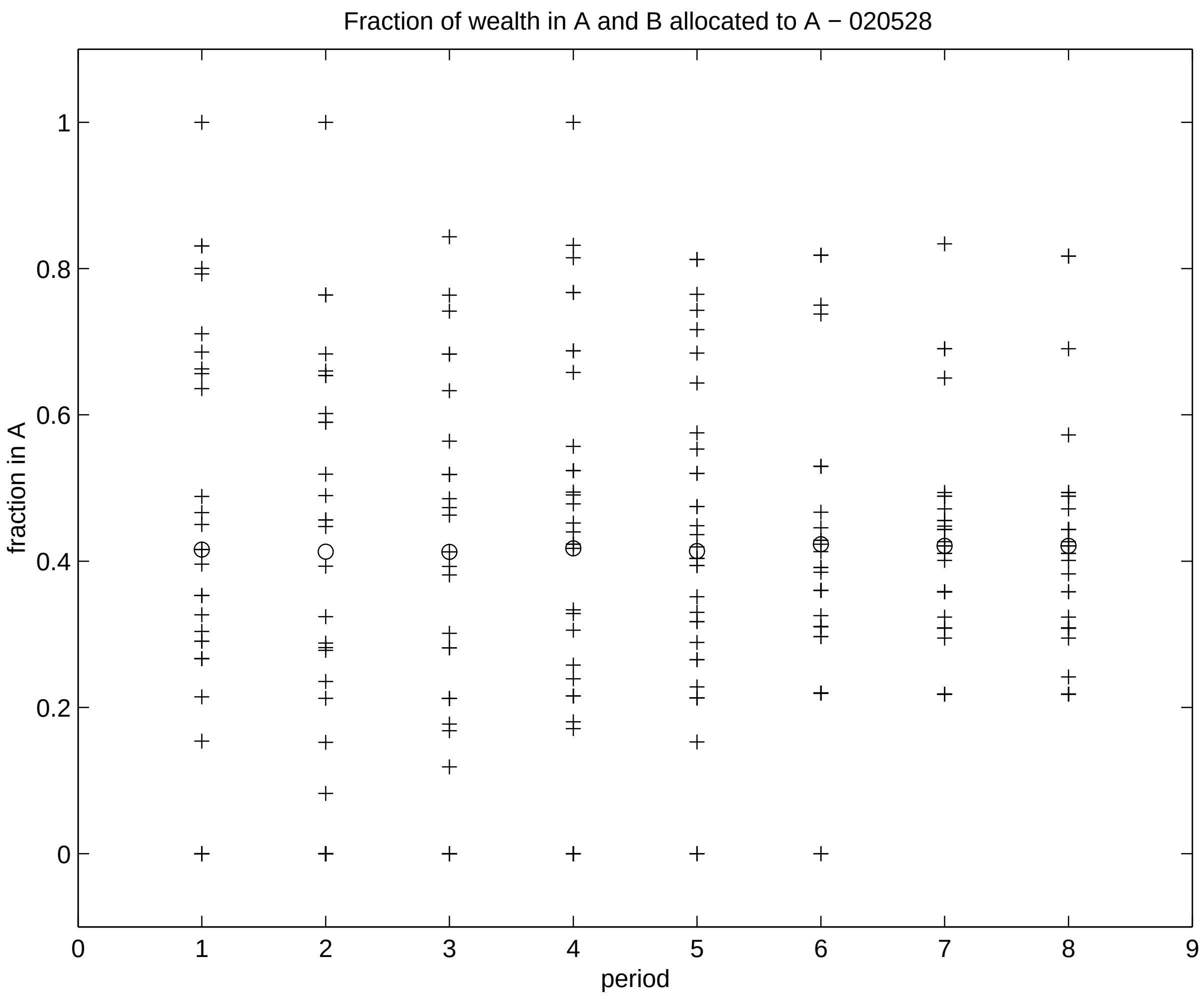

Figure 4: Evolution of end-of-period holdings of risky securities, experiment 28 May 02. Plotted are: (i) subjects' end-of-period holdings of security A as a proportion of all riski-exposed wealth (+), (ii) the end-of-period weight of A in the market portfolio (o).

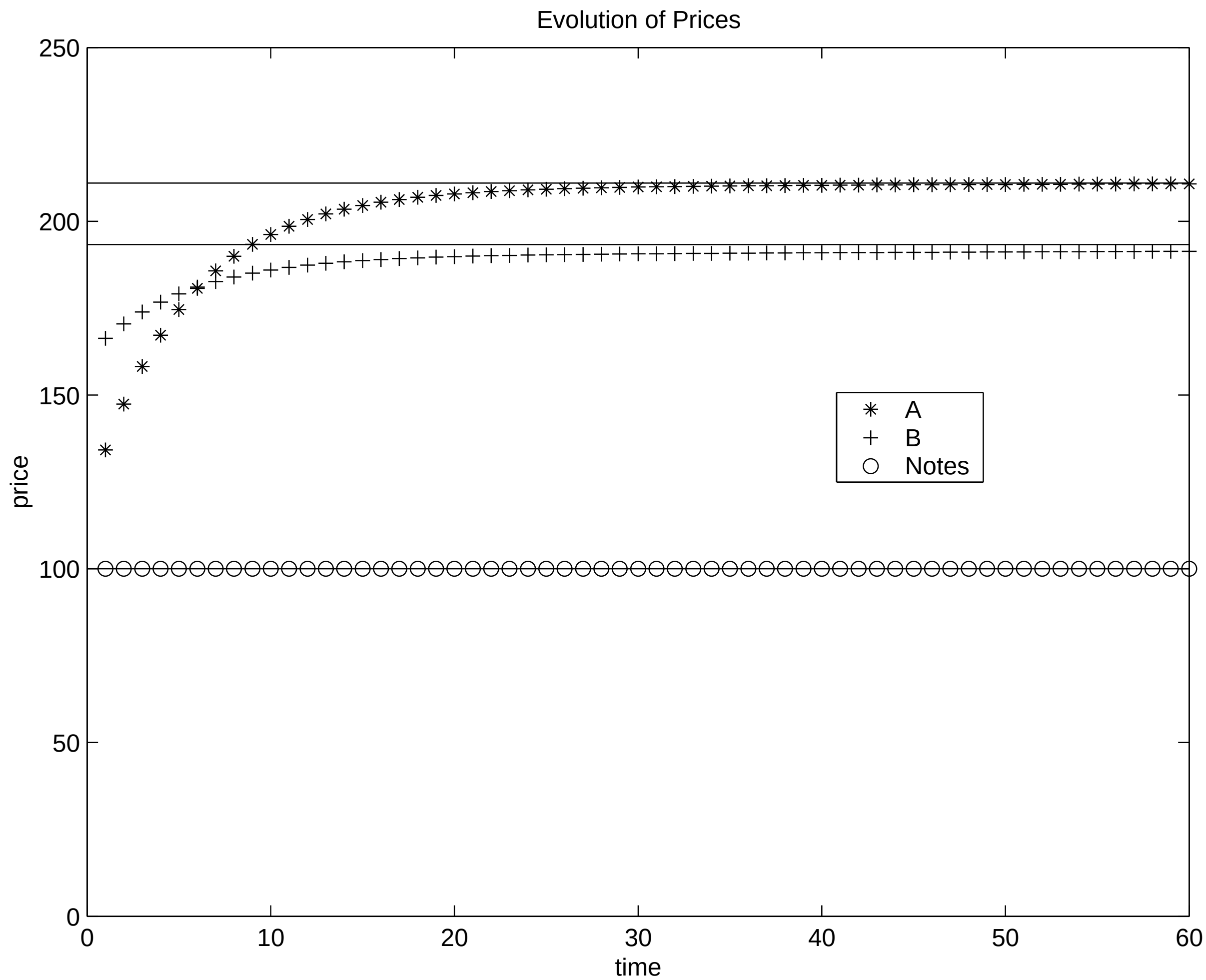

Figure 5: Theoretical evolution of prices in hypothetical markets with the same parameters as experiment 28 May 02, except for initial allocations (all agents start with the market portfolio). Horizontal lines denote equilibrium prices.

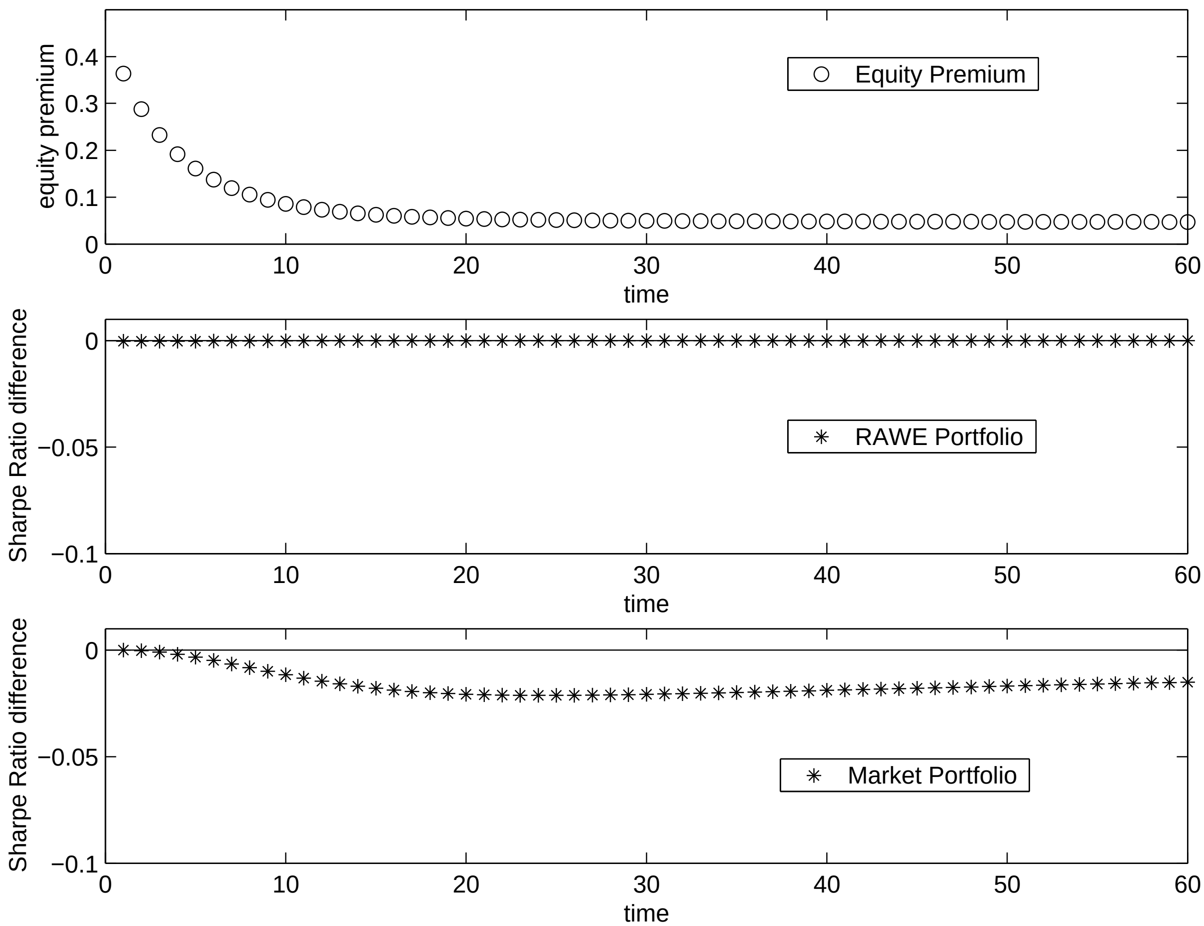

Figure 6: Theoretical evolution of the equity premium (top panel), the distance of the RAWE portfolio (middle panel) and the market portfolio (bottom panel) from the mean-variance efficient frontier in hypothetical markets with the same parameters as experiment 28 May 02, except for initial allocations (all agents start with the market portfolio). Distance of a portfolio from the mean-variance efficient frontier is measured by difference between the portfolio's Sharpe ratio and the maximal Sharpe ratio.

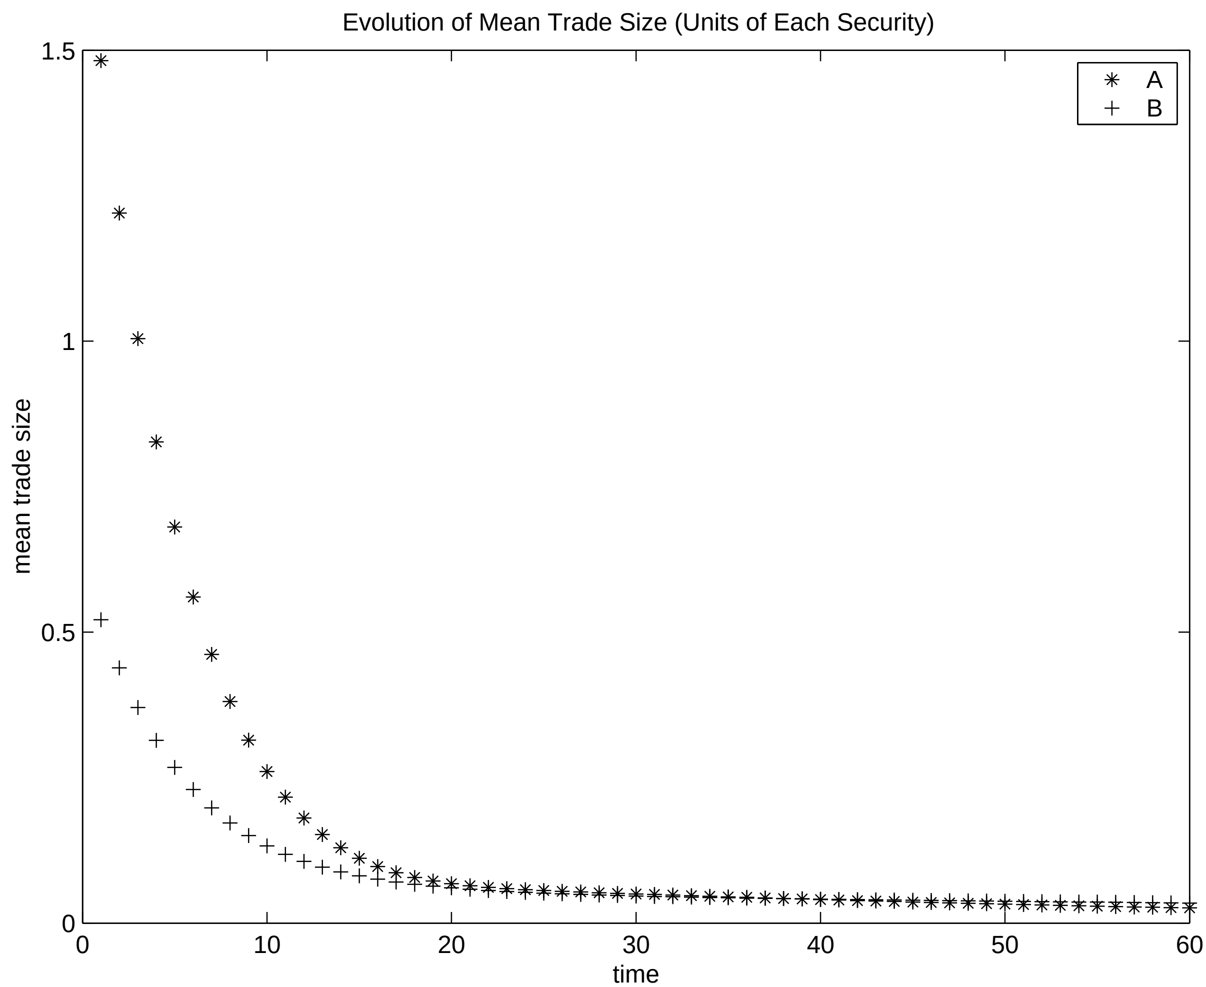

Figure 7: Theoretical evolution of the mean trade size (number of units per agent) in hypothetical markets with the same parameters as experiment 28 May 02, except for initial allocations (all agents start with the market portfolio).

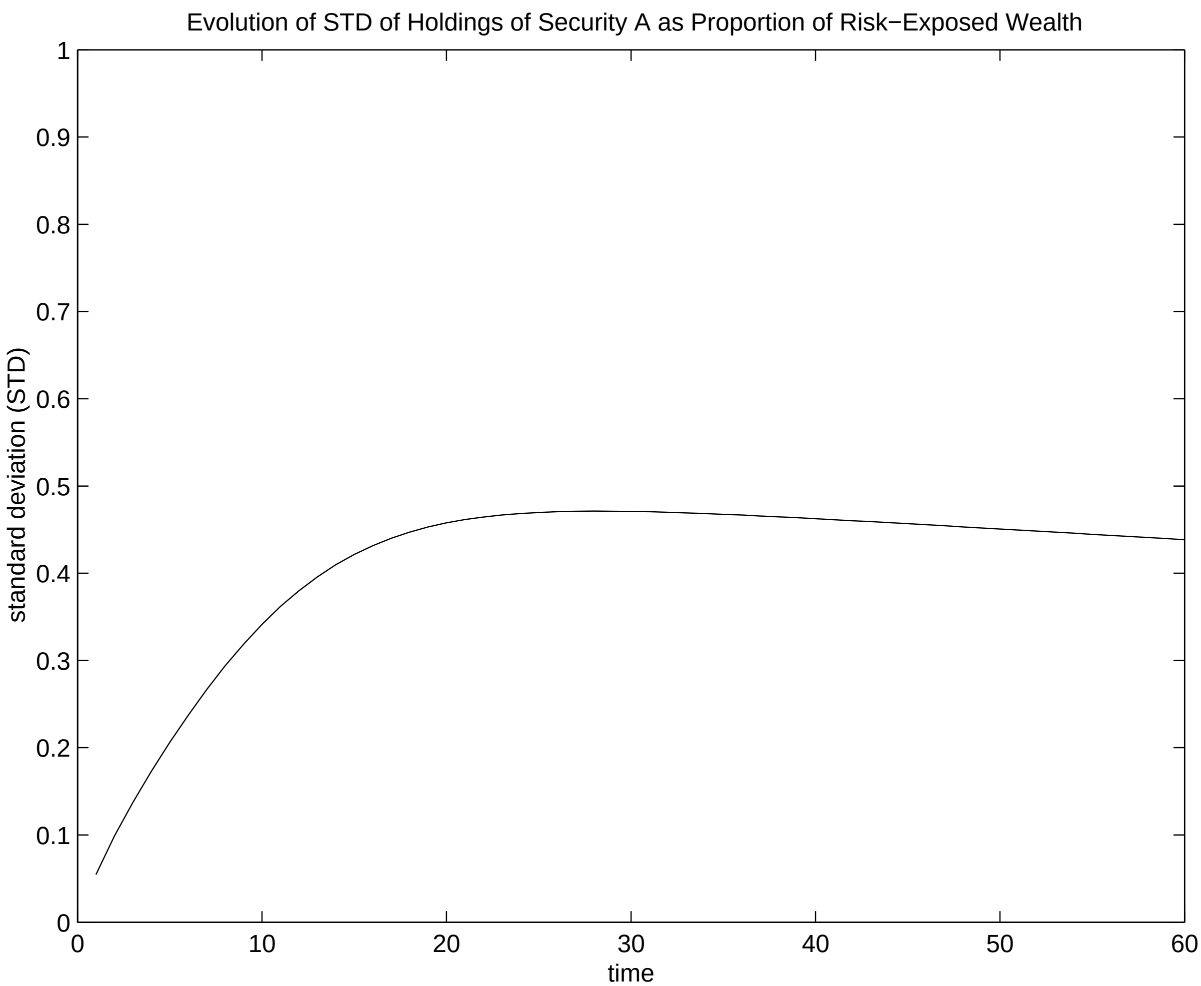

Figure 8: Theoretical evolution of the standard deviation of risky securities holdings across agents in hypothetical markets with the same parameters as experiment 28 May 02, except for initial allocations (all agents start with the market portfolio). Risky securities holdings are measured by the proportion of risk-exposed wealth that agents put in security A.