ELSEVIER

European Economic Review 46 (2002) 1671 – 1695

EUROPEAN ECONOMIC REVIEW

www.elsevier.com/locate/econbase

Inducing liquidity in thin financial markets through combined-value trading mechanisms

Peter Bossaerts[a,b,*], Leslie Fine[c], John Ledyard[a]

[a] California Institute of Technology, Division of Humanities & Social Sciences, 228-77 Caltech, Pasadena, CA 91125, USA
[b] Center for Economic Policy Research, 90–98 Goswell Road, London EC1V 7RR, UK
[c] Hewlett Packard, 3000 Hanover Street, Palo Alto, CA 94304–1185, USA

Abstract

Asset pricing theory hypothesizes that investors are only interested in portfolios; individual securities are evaluated only in terms of their contribution to portfolio risk and return. Yet, standard financial market design is that of parallel, unconnected markets, whereby investors cannot submit orders in one market conditional on events in others. When markets are thin, this exposes them to substantial execution risk. Fear of ending up with unbalanced portfolios after trading may even keep investors from submitting orders, further eroding liquidity and the ability of markets to equilibrate. The suggested solution is a portfolio trading mechanism referred to as combined-value trading (CVT). Investors are allowed to submit orders for packages of securities and the system matches trades and computes prices by optimally combining portfolio orders in an open book. We study the performance of the CVT mechanism experimentally and compare it to the performance of parallel, unconnected double auctions in experiments with similar parametrization and either a similar number of subjects or substantially thicker markets. We present evidence that our portfolio trading mechanism facilitates equilibration to the extent that the thicker markets do. Inspection of order submission and trade activity reveals that subjects manage to exploit the direct linkages between markets enabled by the CVT system.

JEL classification: C92; D44; D51; G12

Keywords: Liquidity; Portfolio trading; Financial markets; Experimental finance

* California Institute of Technology, Division of Humanities and Soc. Sc., 228-77, Caltech, Pasadena, CA 91125, USA. Tel.: +1-626-395-4028; fax: +1-626-405-9841.

E-mail address: pbs@rioja.caltech.edu (P. Bossaerts).

0014-2921/02/$ - see front matter
PII: S0014-2921(02)00240-4

1. Introduction

Asset pricing theory builds on portfolio analysis, i.e., the study of optimal combinations of securities. The key idea is that investors ought not to hold securities singly, but instead buy portfolios, exploiting the less than full cross-sectional dependence in the returns to reduce the risk (measured as variance, skewness, etc.) without affecting expected return. Asset pricing theory assumes that investors understand these principles. Hence, investors' demands in the marketplace follow directly from portfolio optimization. In equilibrium, prices will reflect these demands and highly desirable portfolios will be more expensive than other portfolios. Which implies that the pricing of individual securities cannot be understood without reference to the pricing of other securities or how they fit in portfolios.

The central position of portfolios in the theory of asset pricing contrasts markedly with the way most financial markets are organized. One would expect that these had become vehicles for the exchange of portfolios. Instead, they are mostly set up as parallel markets where securities are traded singly, without the possibility to condition one's orders on interdependent orders for other securities. Even more extremely, the persons who oversee and arrange transactions often make a market in only a few securities (e.g., the specialists on the NYSE). Except for a minority of heavily traded securities, markets are generally too thin and expose investors to a serious execution risk: Desired trades either will not take place, or will happen but at unexpected prices, potentially upsetting the optimality of the resulting portfolio. Bossaerts (1999) provides some estimates of this execution risk when submitting orders at the open of the NYSE. [1]

But it may be that the theory's presumption that investors should focus on portfolios and not individual securities is just wrong. Investors may not care about portfolios and instead have preferences that rank particular securities above any portfolio of other securities one can imagine. After all, investors do not have to have state-separable utility functions to be rational, whereas portfolio theory invariably uses such utility functions to claim superiority of portfolios over single securities. (If one restricts the definition of rationality to Savage's axioms, portfolios are generally superior; but Savage's axioms do not exhaust what rationality could possibly mean, and investors are known to violate these axioms.) If investors prefer to hold securities singly, the fact that most markets are organized as parallel markets in individual securities would not come as a surprise.

Still, there are a few cases where portfolios are directly traded in the marketplace. For instance, the AMEX has recently introduced markets in indices. The introduction has been motivated by portfolio theory. In fact, the nature of the indices (marketwide, value-weighted portfolios) suggests that the architects of these index markets not only

[1] Market microstructure theorists have also suggested that it may enhance liquidity to switch to a portfolio trading mechanism. See, e.g., Wohl (1997) and Wohl and Kandel (1997). The conclusion is based on models with asymmetric information, something we will not be concerned with. With asymmetric information, however, trade in standard portfolios (e.g., marketwide. value-weighted portfolios) is not always liquidity-enhancing. See, e.g., Cespa (2001).

subscribe to portfolio theory, but also to the particular versions of the asset pricing theory that it generates, because the optimality of the exchange-traded indices can only be justified in particular equilibria (e.g., the CAPM equilibrium).

Even if one subscribes to the theory that investors do care only about portfolios, and to particular equilibria that such preferences induce, it is still not clear that it is optimal to trade only the marketwide, value-weighted indices of the AMEX. The optimality depends on investors' endowments, about which little is known. For instance, Athanasoulis and Shiller (2000) provides an example of an economy where it is optimal to trade portfolios that are *orthogonal* to the market portfolio (the value-weighted portfolio of all securities in the economy), even if the CAPM equilibrium obtains. In general, one can expect the nature of the portfolio that individuals desire to trade to change with the nature of the individuals' endowments. Different investors want to trade different portfolios.

This calls for imaginative trading mechanisms, where orders for heterogeneous portfolios are matched. Some one-sided markets in packages are being considered outside the realm of financial markets (e.g., wireless telephony spectrum auctions). They are meant to avert the losses stemming from overbidding when faced with high demand in a market for one of the goods in the desired package.[2] Our purpose is to design and study a two-sided market in heterogeneous portfolios of financial securities. This is an extension of the markets for pollution rights (AQMD) to financial markets. The mechanism has also been proposed for bond markets.[3] In two-sided markets, the aim of portfolio trading is slightly different from that in one-sided markets, namely, to accommodate swaps of packages.

The goal of this article is to propose such a two-sided trading mechanism and to gauge its performance in experiments. The mechanism, referred to as combined-value trading (CVT), is designed to cross heterogeneous portfolio orders in an intermittent call market with an open (portfolio) book. The crossing is accomplished by a scale-back procedure that is reminiscent of the partial order filling in standard, one-security markets. The second novelty of our trading mechanism is pricing. Markets need a clear, easily interpretable signal that reflects excess demand (price increases) or excess supply (price decreases). Constrained, mixed linear-integer programming is used to determine prices and trades. The constraints are suggested by economic theory.

To gauge the performance of the proposed CVT mechanism, we measure the frequency with which investors exploit the ability to submit orders for portfolios (as opposed to individual securities). More importantly, we provide an indirect measure of the welfare that the mechanism manages to create. If markets are complete, the resulting competitive equilibrium is Pareto optimal, and hence, the resulting allocations maximize welfare. This means that we can indirectly measure welfare improvement in terms of the distance from equilibrium. This requires, however, that we know what equilibrium to look for. Recent experiments on financial markets provide an indication.

[2] See Bykowsky et al. (2000).

[3] See Polk and Schulman (2000).

In these experiments, risk is relatively small, which justifies the use of quadratic approximations to investors' actual preferences. Because of this, the ensuing asset market equilibrium is given by the capital asset pricing model (CAPM). That is, prices of securities will be such that the market portfolio (value-weighted index of all risky securities) is mean – variance efficient (provides maximum expected return given its volatility). Even if investors' preferences are not of the expected-utility type, in which case the quadratic approximation may leave a large error, CAPM pricing may still obtain. Indeed, Bossaerts et al. (2001) demonstrate that CAPM pricing will emerge on average as long as the difference between investors' actual security demands and the optimal demand with quadratic utility is mean zero (across investors). Of course, the strong prediction that CAPM makes about equilibrium allocations (all investors hold the same portfolio of risky securities, namely, the market portfolio) will not obtain anymore.

In Bossaerts et al. (2001), these equilibrium predictions are shown to emerge in financial markets experiments with a sizeable number of subjects (relative to traditional market experiments in economics). In particular, the market portfolio tends to become mean – variance efficient, and if deviations occur (e.g., at the beginning of a new period), the market portfolio climbs back to the mean – variance frontier. While allocations are not as predicted by the traditional CAPM (subjects generally buy portfolios that differ substantially from the market portfolio), structural econometric tests reveal that CAPM pricing obtained precisely because the error that the CAPM makes in predicting subjects' demands averages out in cross-section. See Bossaerts et al. (2002).

At the same time, experiments with a number of subjects that is more common in experimental economics (less than 20) produce very different results. While the market portfolio tends to climb towards the mean – variance frontier, Bossaerts and Plott (2002) discovered that the convergence process halts well before it reaches the frontier. The obvious conjecture is that thinness of the market made it difficult for subjects to execute their desired portfolio reallocations. In the face of execution risk, they stop trading well before reaching their unconstrained optimal portfolios, thereby making it impossible for market prices to move all the way to CAPM.

The contrast between the thick-market (up to 63 subjects) and thin-market experiments is all the more telling, because the setting was in many respect the same. This article reports on experiments which replicate the thin-market experiments, but where the CVT mechanism is used. The question is whether CVT causes prices to behave as if one were in the thick-market experiments. That is, is CVT with a few subjects able to move prices as close to CAPM on average as did experiments with up to 63 subjects? [4]

[4] The earlier thick-market and thin-market experiments do not constitute controls in the strict textbook sense of the term. That is, more than one control variable is different. But the same was true across thick-market experiments, none of which was (and could never have been) an exact replication of others (the subject cohort, the exact payoff matrix, the market portfolio, experience, transparency of the book, speed of the system, length of the periods, location, etc.). The earlier experiments demonstrated, however, that the dimensions in which the parametrization is different – except for the number of subjects – have no discernible effect on the pricing. The appendix summarizes the similarities and differences across experiments.

The remainder of the article is organized as follows. The next section describes the CVT mechanism. Section 3 provides details of the experimental setup. Section 4 discusses the experimental results. Section 5 concludes.

2. The CVT mechanism

We envisage a market for portfolios of K securities where trades are settled in cash. That is, to purchase portfolios, subjects pay in cash. When selling portfolios, subjects receive cash.

In the CVT mechanism, subjects submit orders to buy portfolios of securities. A portfolio is a vector q of quantities of each of K securities. The kth element of q will be denoted q^k. Elements of q need not be positive. If $q^k < 0$ for some k, then the subject sells security k as part of the portfolio. If $q^k < 0$ for some k and $q^l > 0$ for some $l \neq k$, then the subject signals his willingness to swap security k for security l.

Along with the vector q, a (cash) price b is submitted. b is the maximum amount in cash that the subject is willing to pay. If $b < 0$, then the subject wants to receive at least $|b|$ to purchase the portfolio q. It is more convenient to interpret such a case as an offer to sell the portfolio $-q$ for at least $|b|$.

One would expect subjects to be more familiar with a trading system where each asset is priced separately, rather than one where they are asked to assign prices to portfolios. To accommodate this, the CVT system asks subjects to submit a price vector p along with the quantity vector q. The CVT system will then calculate b simply as the (dot) product of p and q. We will discuss later how the CVT system uses the price vectors p that subjects submit along with the quantity vectors q in order to resolve nonuniqueness when computing the equilibrium price vector p^* on which cash payments and receipts will be based.

Just like a standard single-asset trading system, CVT allows for partial order fills. In other words, it is possible that only a fraction f of the order is executed. In that case, the subject receives (if a purchase) or delivers (if a sale) the portfolio fq, and pays fb (if a purchase) or receives $-fb$ (if a sale). Unlike in a standard single-asset trading system, however, CVT allows the subject to specify a minimum fill fraction F. In that case, the actual fill fraction f will be at least F (and at most 1), if the order is filled at all. This additional flexibility does complicate the order matching and computation of prices somewhat, as will be explained later on. We will momentarily disregard this complication and analyze the case where subjects are not given additional flexibility, i.e., $F = 0$.

In the experiments, CVT is implemented as an intermittent call market with an electronic open book. That is, subjects submit orders electronically (more than one per subject if so desired), which are all displayed in a book. After a pre-determined time, order submission is stopped, and the system attempts to match the orders in the book and compute prices from which payments are determined. We will discuss shortly how order matching and price computation are done. A new round starts, during which subjects can submit new orders. To facilitate new order submission, the system allows subjects to re-new orders that failed to execute in the previous round, with one click of the mouse.

Order matching is determined as follows. With I bids in the book, the system solves the following linear programming problem:

$$\begin{aligned} &\max \sum_{i=1}^{I} f_i b_i \\ &\text{s.t.} \sum_{i=1}^{I} f_i q_i^k \leqslant 0, \quad k = 1,\ldots,K, \\ &\qquad f_i \geqslant 0, \qquad i = 1,\ldots,I, \\ &\qquad f_i \leqslant 1, \qquad i = 1,\ldots,I. \end{aligned} \tag{1}$$

The idea is to maximize surplus (the criterion function), i.e., the difference between the sum of the bids of offers to buy ($b_i \geqslant 0$) that will be filled ($f_i > 0$) and the asks of offers to sell ($b_i < 0$) that will be filled ($f_i > 0$). The surplus is maximized subject to a resource constraint.[5] Note that we do not impose strict market clearing. According to the first constraint, there may be a net surplus. If that happens, the system takes the surplus and offers it for sale in the next round (if there is any).

Once orders are matched in this way, prices, and hence, payments are determined, as follows. The K-dimensional price vector p^* is chosen such that Walras' law obtains

$$p^* \sum_{i=1}^{I} f_i q_i = 0. \tag{2}$$

In addition, the payments/receipts for *filled* orders are never more/less than bid/asked:

$$p^* q_i \leqslant b_i \quad \text{for } i : f_i > 0. \tag{3}$$

Finally, the payments/receipts for *unfilled* orders are always more/less than bid/asked:

$$p^* q_i \geqslant b_i \quad \text{for } i : f_i = 0. \tag{4}$$

Given sufficient regularity in subjects' preferences (namely, the conditions required for the existence of a competitive equilibrium), a solution p^* to the above (in)equalities will always exist. The solution will not be unique because, among other things, we have not assigned a numéraire. If there is a risk-free security that is equivalent to cash, we could set its price equal to 1, to reflect the fact that trades are settled in cash, and hence, prices are expressed in monetary units. In our implementation of the CVT system, we did not do so, opting instead for a rule that is applicable even when there is no risk-free security that is equivalent to cash. Our rule picks the solution p^* with reference to the price vectors p that subjects submitted along with their portfolio orders

[5] We have been interpreting orders with $b_i < 0$ as offers to *sell* $-q$. In keeping with this interpretation, the first constraint should have been written as the difference between the offers to buy (those with $b_i \geqslant 0$) and the offers to sell (those with $b_i < 0$). But the result is the same:

$$\sum_{i:b_i \geqslant 0} f_i q_i^k - \sum_{i:b_i<0} f_i(-q_i^k) = \sum_{i=1}^{I} f_i q_i^k.$$

q. In particular, we take p^* that minimizes the average Euclidean distance with the submitted p. This way, p^* was chosen to reflect what subjects thought the individual securities were worth to themselves.

Our pricing rule also resolves an annoying issue when there is excess supply for a security given the orders in the book. These are situations where the resource constraint is nonbinding for some k:

$$\sum_{i=1}^{I} f_i q_i^k < 0.$$

In principle, one could assign a zero price to this security ($p^{k,*} = 0$). Instead, we pick a price that is closest to subjects' valuations revealed in their orders (their p^ks), albeit ensuring that Walras' law obtains, that payments for filled buy orders are never larger than the bid price, that receipts for filled sell orders are never below the asked price, that payments for unfilled buy orders would at least have been as large as the bid price, and that receipts for unfilled sell orders would at most have been as large as the asked price (see expressions (2)–(4) above).[6]

As mentioned above, the actual CVT mechanism is more complex in one dimension: We allow subjects to submit a minimal fill fraction F. This means that the actual fill fraction f is either zero or somewhere between F and 1. It turns the linear programming problem in (1) into a mixed linear-integer program, because the second and third constraints change to

$$f_i \in [F_i, 1] \cup 0. \tag{5}$$

Our allowing minimal fill fractions could cause a serious problem. If subjects do use the option of specifying a minimal fill fraction in their orders, they reveal that their preferences over securities exhibit nonconvexities. It is well known that a competitive equilibrium may not exist if preferences exhibit nonconvexities. The implication is that CVT may not be able to find a price vector p^* that satisfies the restrictions in expressions (2)–(4), because these restrictions characterize a competitive equilibrium. The appendix discusses how CVT goes about resolving the problem of nonexistence of a solution p^*.[7]

3. Experimental design

We conducted a total of seven experiments. Specifics are given in Table 1, where the experiments are indexed by the date of the session.

Subjects were recruited from the Caltech community, primarily undergraduates and a few graduate students from the natural sciences and engineering. Because of the

[6] By ensuring that Walras' law holds – see (2) – the CVT system always receives enough money to pay for any excess supply it is buying from subjects.

[7] Nonexistence of equilibrium is a rare event in the experiments, though. Of the 442 rounds total, there were only eight instances where accommodation of inflexibilities was necessary because of equilibrium nonexistence.

Table 1
List of CVT experiments and parameters

Date	Subject number[a]	Signup reward (franc)	Endowments			Cash (franc)	Loan (franc)	Exchange rate ($/franc)
			A	B	Notes			
11/1/99	3	430	9	1	0	400	2500	0.023
	3	430	1	9	0	400	2400	0.023
11/4/99	7	430	9	1	0	400	2500	0.023
	7	430	1	9	0	400	2400	0.023
11/11/99	4	430	9	1	0	400	2500	0.023
	7	430	1	9	0	400	2400	0.023
11/16/99	9	430	9	1	0	400	2500	0.023
	5	430	1	9	0	400	2400	0.023
11/30/99	10	430	9	1	0	400	2500	0.023
	5	430	1	9	0	400	2400	0.023
12/2/99	6	430	6	3	0	400	2250	0.03
	3	430	3	6	0	400	2200	0.03
	3	430	4	5	0	400	2220	0.03
12/7/99	9	430	9	1	0	400	2500	0.03
	5	430	1	9	0	400	2400	0.03

[a]Subjects differed in terms of initial allocations of risky securities. The number of each type changed across periods, as initial allocations were randomly assigned, and because subjects had to leave as a consequence of bankruptcy (subjects that had negative cumulative earnings two periods in a row were barred from further trading).

complexity of the trading interface and of the intricacies of securities trading, recruiting was limited to those who had taken or were in the midst of taking courses related to finance. Some subjects, particularly in later session, participated in more than one session. The number of subjects varied from a low of six to a high of fourteen. This is typical for market economics experiments, but substantially less than the asset market experiments discussed in Bossaerts et al. (2001), where up to 63 subjects traded in large web-based markets.

Subjects were allocated varying quantities of three securities, denoted A, B, and C, each with a life of one period. At the end of the period, securities paid a liquidating dividend and were then retired. The magnitude of the dividends depended on the random draw of one of three equally likely states, X, Y or Z. The state was drawn after the period was closed, so there was no insider, i.e., no asymmetric information. The payoff table was as follows:

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

Notice that the dividend of A varied dramatically from state to state (with an expected value of 230), the dividend of B varied less and had an expected value of 200, and the dividend of C was constant at 100. All subjects were given this payoff table.

Each experiment consisted of multiple periods of similar trading conditions, varying only by the initial allocations given to subjects at the beginning of each period. In addition to securities, each player was supplied with some francs cash at the beginning of each period. Cash was the medium of exchange and we used an artificial currency (francs) so that there was no need to move beyond integers to express prices. The conversion rate into dollars was pre-announced (see Table 1 for the exchange rates).

Each period consisted of a number of rounds, the first round 3 minutes in length and subsequent rounds 90 seconds long. During a round, subjects submitted orders which are then displayed for all to see in an electronic open book. At the end of each round, order submission was stopped while the allocation algorithm (described in Section 2) solved for all trades and prices. After trades were executed (if any), another round began. The number of rounds in a period varied between experiments. The November experiments had 10 rounds per period. As will be reported in Section 4, it appeared that markets equilibrated (measured both in terms of an equilibrium model and by the decrease in trades) and subjects slowed their activity after 5–8 rounds. Therefore, the December experiments had only seven rounds per period.

At the end of a period, a state was chosen at random, players were paid liquidating dividends according to the payoff table, and the period ended. Afterwards, another period began, with new allocations. In the case of the November experiments, the session consisted of six periods, while the December sessions were eight periods long. Subjects knew about the length of the session they were in.

Note that while we changed the aggregate supply of securities A and B from period to period, security C was always kept in zero net supply. No short sales were permitted in A and B, but players were allowed to short up to 8 units of C.[8] Thus, the ability to buy A and B is not limited by the number of francs on hand. If a subject wishes to expand holdings beyond the bounds indicated by the endowment of francs, she could do so by selling units of C and paying the dividend. A sale of C is equivalent to borrowing an amount equal to the sale price, in exchange for a repayment of 100 francs at the end of the period. To the buyer, the difference between the price paid and the dividend is a risk-free return since the payment is guaranteed. This sale and purchase of asset C determines the risk-free rate simultaneously with the rates of return on the risky securities.

Since it is possible for players to lose money due to unfortunate draws of the state and to then declare bankruptcy, we needed to ensure the integrity of the incentive system. Any subject whose cumulative earnings were negative in two consecutive periods

[8] In the ensuing equilibrium and assuming expected utility preferences, shortsale constraints on risky securities will never be binding, and hence, are irrelevant. Shortsales of risk-free securities may be important, though, as they allow more risk-tolerant subjects to take away risk from more risk-averse subjects. Therefore, we allowed for generous shortselling of risk-free securities.

in an experimental session was asked to leave the experiment. Even if a subject is risk neutral, this incentive system does induce risk-averse behavior up to and including the penultimate period, because bankruptcy is equivalent to foregoing an opportunity to make more money. As will be documented later, however, risk aversion emerges in the pricing of securities from the beginning of each experiment. It always lasted till the end, casting doubt that the observed risk aversion was only induced through the bankruptcy rule.

The initial allocations varied from experiment to experiment, as detailed in Table 1. While players were aware of their own initial allocations, they were not told the allocations of others. Therefore, the composition of the market portfolio (portfolio of all risky securities in the market place) remained unknown. This is an important design consideration. It ensures that subjects could not use the CAPM (the model that will be used to evaluate whether markets equilibrated) to deliberately set prices, thereby artificially generating the outcomes that we were looking for.

Before assembling in the Caltech Social Science Experimental Laboratory, all subjects were given a URL with the instructions. At the end of each page of instructions is a short quiz, which had to be correctly completed before moving on to the next page. Once all pages were completed, subjects were given a URL for a practice experiment. Subjects were not allowed to participate in the 3 hours in-lab experiment if they did not enter at least five practice bids in the practice experiment. Each subject who completed these tasks and arrived at the lab on-time received a \$10 bonus (which was fully exposed to risk during the experiment, though). On average and excluding bankruptcies, subjects made roughly \$60 per experiment (experiments usually lasted about $2\frac{1}{2}$ hours).

The appendix provides details of the two sets of control experiments: Thick-markets experiments and thin-markets experiments, both with a parallel system of markets where explicit coordination of orders across markets is infeasible. The control experiments were discussed in more detail in Bossaerts and Plott (2001, 2002), Bossaerts et al. (2001, 2002). The control experiments do not constitute controls in the strict textbook sense of the term. That is, more than one control variable is different. But the same was true across all experiments within a category (CVT, thin-market, thick-market), none of which was (and could never have been) an exact replication of others (the subject cohort, the exact payoff matrix, the market portfolio, experience, transparency of the book, speed of the system, length of the periods, location, etc.). The experiments demonstrated, however, that the dimensions in which the parametrization is different – except for the number of subjects – hardly affects the outcomes.

4. Results

We will gauge the success of the CVT trading mechanism in four ways: (i) informal evidence of equilibration through reduction in volume, (ii) formal tests of equilibration (relative to CAPM) in experiments with CVT, (iii) comparison of equilibration (relative to CAPM) between experiments with CVT and experiments with parallel markets, (iv) relative usage of packages (portfolios) in order submission.

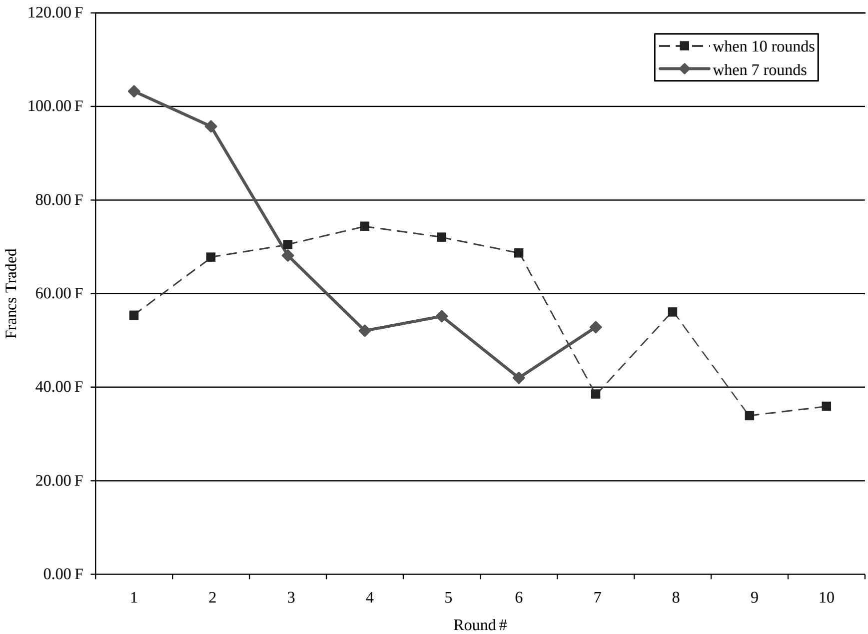

Fig. 1. Per capita transaction volume (in francs), by round, averaged across periods and experiments.

4.1. Informal evidence of equilibration

One crude way to ascertain that markets equilibrated is to observe patterns in volume. An indication of equilibration is that all gains from trade are exhausted and volume dies down.

Fig. 1 depicts the evolution of volume (in francs traded per subject) over time. In the ten-round experiments (11/99 experiments), volume remains even over the first six rounds, and declines subsequently. This is a simple indication that most gains from trade are exploited by round six. In the seven-round experiments (12/99 experiments), volume declines already after two rounds. Apparently, the reduction in number of rounds forces subjects to reveal their desired trades earlier, and many gains from trade are thereby exploited faster.

Of course, the above is only casual evidence that markets equilibrated. We would like to have more formal measures of equilibration, and if possible, formal tests that prices indeed move in the direction of equilibrium. These are introduced next.

4.2. Formal tests of equilibration

To generate formal measures of equilibration, we need an asset pricing model. The limited size of the stakes in a typical experiment justifies approximating subjects' preferences towards risk by mean – variance utility functions. The resulting equilibrium is given by the capital asset pricing model (CAPM). Even if mean – variance preference theory does not accurately describe subjects' actual attitudes towards risk, Bossaerts et al. (2001) demonstrate that CAPM pricing will still obtain in expectation as long as mean – variance preference theory adequately describes individual portfolio demands *on average*. That possibility, however, precludes our using end-of-period holdings as a measure of distance from equilibrium. [9]

CAPM predicts that equilibrium prices will re-arrange such that the market portfolio (i.e., the aggregate supply of risky securities) is optimal for mean – variance preferences. That is, the market portfolio generates maximum mean return for its volatility. This prediction is independent of agents' levels of risk aversion, which is fortunate, because these cannot readily be measured. Moreover, risk aversion may change during the course of the experiment, as in, e.g., the previously mentioned case where subjects are risk-neutral, but our bankruptcy rule nevertheless induces risk aversion up to the penultimate period.

Consequently, to determine whether experimental markets have equilibrated, one can compare the reward-to-risk trade-off of the market portfolio against the maximum possible trade-off available at market prices. This trade-off is usually referred to as *Sharpe ratio*, to be defined as follows. Let R_{Ft} denote the return on a risk-free security at time t; let R_{mt} be the return on the market portfolio and let σ_{mt} denote its volatility. [10] The Sharpe ratio of the market portfolio is then

$$\frac{\mathrm{E}[R_{mt} - R_{Ft}]}{\sigma_{mt}}.$$

The maximal Sharpe ratio is the maximum possible ratio of mean return in excess of the risk-free rate over volatility (when a risk-free security exists, the maximal Sharpe ratio is constant for all levels of volatility). [11]

At any moment in our experiments, we measure distance from equilibrium by computing the difference between the market Sharpe ratio and the maximal Sharpe ratio.

[9] Even in thick-market experiments, where prices predominantly converge to a CAPM equilibrium configuration, end-of-period allocations invariably deviate significantly from those predicted by the CAPM (where every agent is expected to hold risky securities in the same proportion). See Bossaerts et al. (2001) for evidence and Bossaerts et al. (2002) for formal statistical tests that CAPM pricing obtain because mean – variance preference theory only describes individual portfolio choice on average. In previous versions of this article, end-of-period holdings in the experiments with the CVT trading mechanism were reported to deviate likewise from CAPM allocations.

[10] Returns are computed as final period payoff divided by last transaction prices as of time t. Expected returns are computed using the announced probabilities of each possible final payoff; variances and covariances are computed analogously.

[11] In our computation of the maximal Sharpe ratio, we did take into account constraints on shortselling of risky and risk-free securities. In particular, optimal portfolios are combinations of the risk-free security and some portfolio of the risky securities usually referred to as the tangency portfolio. In the tangency portfolio, negative weights are not allowed. In addition, the weight on the risk-free security is not allowed to be such that more than 8 risk-free securities are to be sold short. As a result, the mean – variance frontier is not a linear function in mean/standard deviation space and the maximal Sharpe ratio changes with the volatility in its denominator. We always use the maximal Sharpe ratio corresponding to the volatility of the market portfolio.

Markets reach equilibrium when the difference becomes zero. This measure was first successfully employed in an experimental setting in Bossaerts and Plott (2001).

Before reporting the empirical results, it should be emphasized that CAPM equilibration is not a foregone conclusion even in our relatively simple experimental environment. This is because endowments, and hence, aggregate risk (the dividend of the market portfolio), were not common knowledge. To avoid that subjects could learn the aggregate risk over time, we randomly changed individual initial allocations across periods within an experiment, as well as across experiments. [12] Absent information on aggregate risk, one cannot derive CAPM equilibrium prices. Consequently, if we observe equilibration, it is only because of the economic forces that are at work, and not because of subjects' deliberate appeal to the CAPM in determining offer prices.

In addition, the CVT clearing mechanism *in no way* uses CAPM to determine transaction prices. It strictly computes those prices on the basis of the submitted orders, maximizing gains from trade. See Section 2. CAPM price configurations will emerge only if mean-variance demands are reflected in submitted orders.

Fig. 2 shows plots of the evolution of Sharpe ratio differences for our seven experiments. Sharpe ratio differences are computed after every round within a period, on the basis of the transaction prices computed by the CVT mechanism. With few exceptions, the (absolute) Sharpe ratio differences are less than 0.15. There is a tendency for the differences to be wider in earlier rounds of a period, but to narrow subsequently. This indicates that markets move towards CAPM equilibrium. In later rounds differences sometimes widen, which would appear to imply that the market moves off its equilibrium. But volume decreases as well, making prices more sensitive to the few orders that are executed, or in the absence of trade, to the offer prices. Market prices are then more likely to move away from CAPM predictions.

We also reported earlier that reducing the number of rounds forces trading (transaction volume) to occur earlier in a period. Hence, if our previous reasoning is correct, and reduction of trading volume indicates equilibration (interpreted as exhaustion of gains from trade), we expect to still see the same pricing arbitrariness in later rounds of experiments with fewer rounds. This is confirmed in the last two frames of Fig. 2, where the results are plotted for the December experiments, which featured seven instead of ten rounds per period.

Subjects sometimes submitted orders that merely attempted to exploit mistakes of others. This was especially apparent in the orders for the risk-free asset (security C) during later rounds of a period. While the theoretical equilibrium price is 100, subjects sometimes bid a low price for security C, and received it when others inadvertently submitted a low ask. Such speculative offers could have a major impact on the prices computed by the CVT mechanism (which uses offered prices as reference if there is excess supply – see Section 2). Because of this, the price of C was set equal to its

[12] In the thick-market experiments reported in Bossaerts et al. (2001), individual initial allocations did not change from one period to another. Even so, subjects could still not readily determine the distribution of endowments, because the trading environment was far less transparent (trade took place over the internet, and hence, subjects could not see each other; subjects were not informed about (the sizeable number of) other subjects; etc.).

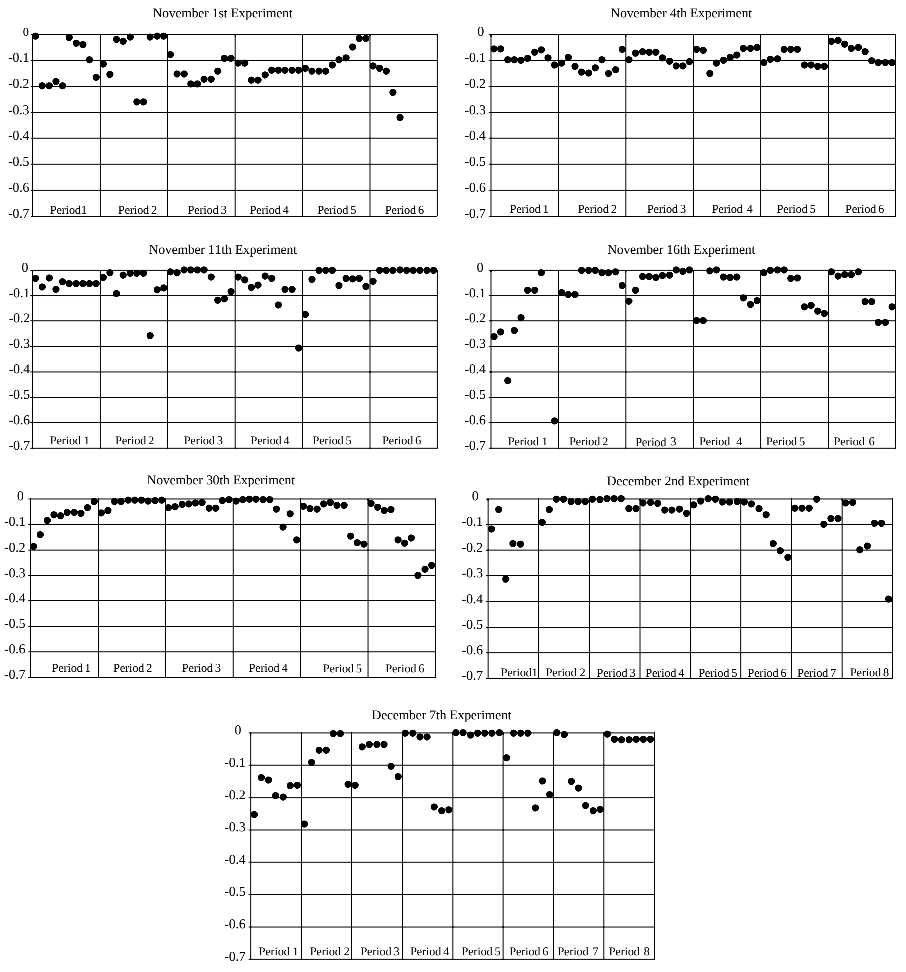

Fig. 2. Evolution of the distance from CAPM pricing equilibrium (difference between the Sharpe ratio of the market portfolio and the maximum Sharpe ratio).

no-arbitrage value (100) before computing the Sharpe ratios differences that are plotted in Fig. 2.

So far, we have relied on graphical evidence of equilibration. We should ask whether our results have not emerged by chance. Even if they are not determined by any economics at all, prices may accidentally move towards CAPM predictions, i.e., the market portfolio may accidentally become mean – variance optimal. To reject this possibility, we take an information-theoretic approach. Absent formal asset pricing theory, all price changes should be complete surprises. That is, price changes constitute a martingale difference sequence. Against this null hypothesis, we determine to what extent the CAPM provides information to predict price changes. If the market portfolio is mean – variance suboptimal, prices should change in predictable ways to move the market portfolio closer to the mean – variance efficient frontier.

To determine the significance level of predictable reversion to the mean – variance frontier, we follow a bootstrapping approach. Specifically, let $\Delta_{M,t}$ denote the distance between the Sharpe ratio of the market and the maximum Sharpe ratio, the subscript t denoting time (period and round). Consider the projection of the change in $\Delta_{M,t}$ onto $\Delta_{M,t-1}$:

$$\Delta_{M,t} - \Delta_{M,t-1} = \kappa\Delta_{M,t-1} + \varepsilon_t \tag{6}$$

where κ is such that ε_t is uncorrelated with $\Delta_{M,t-1}$. CAPM implies $\Delta_{M,t} = 0$; convergence to CAPM pricing implies $\kappa < 0$. We then determine the distribution of the least-squares estimates of κ under the null hypothesis of a martingale difference, by randomly drawing from (bootstrapping) the empirical joint distribution of price changes. The null hypothesis is rejected in favor of CAPM if the least-squares estimate of κ is beyond a critical value in the left tail of the ensuing distribution.

For each experiment, we estimated κ using OLS. We determined 5% and 10% critical values under the null hypothesis by bootstrapping from the empirical joint distribution of price changes (we generated 200 price series of the same length as the sample used to estimate κ). We bootstrapped the mean-corrected empirical distribution, in order to stay with the null hypothesis of a martingale difference sequence. Our testing procedure is a variation of *indirect inference*, first suggested in Gouriéroux et al. (1993): We summarize the data in terms of a simple statistical model (in our case, a least-squares projection) and determine the distribution of the estimator by simulating the variables entering the statistical model. Note that the model in (6) may be mis-specified even under the null that prices are a martingale. Because $\Delta_{M,t}$ is a nonlinear transformation of prices, changes in $\Delta_{M,t}$ may be predictable (κ is different from zero), and the error term may have a nonzero mean, and even exhibit heteroscedasticity while price changes are homoscedastic. The true dynamics of $\Delta_{M,t}$ will in general be nonlinear. This is the main reason why we determine the properties of the OLS estimator of κ on the basis of a bootstrap, and not on the basis of the (generally faulty) assumption that (6) is well-specified.[13] The bootstrap will also accommodate skewness and kurtosis in price changes without having to model these properties explicitly.

Notice that our formal statistical test of equilibration differs substantially from those used in the analysis of field data. Econometric techniques developed to test asset pricing theory on field data are simply not appropriate for the analysis of experimental data. In experiments, one controls crucial parameters that are unknown to the empiricist working with field data. Most importantly, the field empiricist does not know the parameters of the payoff distribution. The usual formal statistical tests build on this.

[13] In indirect inference, one usually simulates off a theoretical distribution. Instead, we bootstrap the empirical distribution. Indirect inference is also used in Gallant and Tauchen (1996), but instead of matching an arbitrary statistical model as Gouriéroux et al. (1993) and we do, the scores of the theoretical likelihood function are matched in Gallant and Tauchen (1996).

Table 2
Predictability of CAPM equilibrium pricing

Experiment	Attraction coefficient κ		
	Estimate[a]	Critical value	
		5%	10%
11/01/99	−0.0893*	− 0.1086	− 0.0789
11/04/99	−0.0400	− 0.1335	− 0.1028
11/11/99	−0.3900**	− 0.1153	− 0.0662
11/16/99	−0.3700**	− 0.0978	− 0.0702
11/30/99	−0.0300	− 0.0942	− 0.0786
12/02/99	−0.2900**	− 0.1325	− 0.0818
12/07/99	−0.2800**	− 0.1001	− 0.0812

[a]Based on 200 bootstrapped samples of the same size as used to estimate κ.

*Significant at the 10% level. ** Significant at the 5% level. Overall significance: Less than 1%.

Implementing such tests in an experimental setting would amount to econometric schizophrenia, because the experimentalist does know the parameters of the payoff distribution.

Table 2 reports the results. The null hypothesis is rejected in one experiment at the 10% level, in four at the 5% level. There is less than a 1% chance to see four or more rejections at the 5% level in seven trials. Overall, therefore, we reject the null of unpredictable price changes in favor of CAPM at a significance level less than 1%. To put this differently, the statistical tests indicate that there is only a tiny probability to accidentally obtain the dynamics of Fig. 2 if price changes were truly unpredictable.

At one point, Fig. 2 suggests that our test may not be very powerful, rendering further credibility to the rejections. Specifically, in the 11/30/99 experiment, the volatility is so low that, given opening prices that positioned the market portfolio close to the mean – variance efficient frontier, one could easily have moved onto the frontier accidentally.

4.3. Comparison with other experiments

We set out to study how well our CVT mechanism does *relative to* the usual system of parallel, continuous double auctions with the same number of subjects. To gain perspective, Table 3 first reports the averages of the distance between the Sharpe ratio of the market and the maximal Sharpe ratio, for all periods of each experiment. Table 3 also reports standard deviations, in parentheses. Within a period, the distance from CAPM equilibrium price configuration is clearly not independent over time: It displays pronounced trending. See Fig. 2. Therefore, we do not report usual standard errors (the standard deviation divided by the square root of the time series length), which would be a good measure of the estimation error of the sample mean only for independent observations. We report the standard deviations instead.

Table 3
Average distance from CAPM equilibrium pricing: CVT experiments against thin and thick parallel-market experiments

Experiment	Periods[a]							
	1	2	3	4	5	6	7	8
11/01/99[b]	0.11	0.10	0.14	0.14	0.11			
	(0.09)	(0.10)	(0.06)	(0.03)	(0.04)			
11/04/99	0.08	0.13	0.09	0.09	0.09	0.07		
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.04)		
11/11/99	0.05	0.06	0.03	0.06	0.07	0.01		
	(0.01)	(0.08)	(0.05)	(0.03)	(0.10)	(0.02)		
11/16/99	0.27	0.09	0.04	0.07	0.07	0.09		
	(0.27)	(0.18)	(0.04)	(0.08)	(0.07)	(0.09)		
11/30/99	0.08	0.02	0.02	0.02	0.07	0.14		
	(0.05)	(0.02)	(0.01)	(0.04)	(0.06)	(0.10)		
12/02/99	0.32	0.02	0.01	0.03	0.01	0.12	0.05	0.15
	(0.28)	(0.03)	(0.02)	(0.02)	(0.01)	(0.11)	(0.03)	(0.13)
12/07/99	0.18	0.09	0.08	0.11	0.00	0.10	0.13	0.02
	(0.04)	(0.10)	(0.05)	(0.13)	(0.00)	(0.10)	(0.09)	(0.01)
all CVT markets	0.16	0.07	0.06	0.07	0.06	0.09	0.09	0.09
	(0.10)[c]	(0.04)	(0.05)	(0.04)	(0.04)	(0.05)	(0.06)	0.09
Thin markets[d]	0.17	0.10*	0.18**	0.17**	0.14**	0.15**	0.12	0.14
	(0.15)	(0.05)	(0.10)	(0.10)	(0.10)	(0.03)	(0.16)	(0.06)
Thick markets[e]	0.20	0.07	0.06	0.07	0.04	0.08	0.06	0.03
	(0.18)	(0.07)	(0.06)	(0.07)	(0.02)	(0.30)	(0.09)	(0.02)

[a]Average difference between the maximum Sharpe (reward-to-risk) ratio and the Sharpe ratio of the market portfolio. Standard deviations in parentheses.

[b]Combined value trade (CVT) experiments are identified with date.

[c]Standard deviation of averages across experiments.

[d]Median (or one smaller if the median is not equal to a realized outcome) of the average difference between the maximum Sharpe (reward-to-risk) ratio and the Sharpe ratio of the market portfolio, experiments with thin parallel double auction markets (see Bossaerts and Plott, 2002). Corresponding standard deviations in parentheses. An asterisk indicates that the number of average differences in the CVT experiments below the median in the thin-market experiments is significant at the 10% level. Two asterisks indicates significance at the 5% level. Overall significance: less than 1%.

[e]Median (or one smaller if the median is not equal to a realized outcome) of the average difference between the maximum Sharpe (reward-to-risk) ratio and the Sharpe ratio of the market portfolio, experiments with thick parallel open-book markets (see Bossaerts and Plott, 2001). Corresponding standard deviations in parentheses. The number of average differences in the CVT experiments below the median in the thick-market experiments is nowhere significant at the 10% level.

To make formal comparisons with other experiments, Table 3 reports the median average distance for eight experiments with thin, parallel markets as well as six thick, parallel markets experiments. The thin-market experiments are based on similar numbers of subjects as the experiments with the CVT trading mechanism. They are described in Bossaerts and Plott (2002). The thick-market experiments were first reported on in Bossaerts and Plott (2001) and further analyzed in Bossaerts et al. (2001, 2002). As mentioned before, the number of subjects in the latter category of experiments was substantially larger (up to 63 subjects). With the exception of a few minor details, the design of all these experiments (including payoff parameters) was the same. See the appendix for further details.

The median is obtained as follows. We first compute the distance between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio after each transaction in a period. We then average across transactions. It is the median of these averages for the same period across experiments that is reported in Table 3. [14] Corresponding standard deviations (of the average distance across experiments) are displayed in parentheses.

To test whether the CVT mechanism improves upon parallel, continuous double auctions with a similar number of subjects, we proceed as follows. We determine the number of CVT experiments where the average distance from CAPM (across rounds within a period) is below the corresponding period median of the thin, parallel-market experiments. We then compute the significance level of this number, under the null that the average distance in a CVT experiment is equally likely above or below the median of the thin, parallel-market experiments. One asterisk indicates that this number is significantly higher than 50% at the 10% level; two asterisks indicate significance at the 5% level. In most periods, the average distance from CAPM in the CVT experiments is below the median of the thin-market experiments more often than expected by chance. Across periods, the overall significance level is less than 1%, suggesting that the CVT mechanism outperformed the parallel double-auction system of the thin-market experiments.

The significance of this finding is enhanced when one also takes into account that the CVT experiments were tougher in at least one respect: Subjects' endowments were randomly changed across periods, whereas in the thin-market experiment, subjects were given the same initial allocation in all periods. [15]

To gain further perspective, we repeat this exercise for the thick, parallel-market experiments. That is, we determine the number of CVT experiments where the average distance from CAPM (across rounds within a period) is below the corresponding period median of the thick, parallel-market experiments. We then compute the significance level of this number, under the null that the period average distance in a CVT experiment is equally likely to be above or below the median for the thick, parallel-market experiments. The frequency that the average distance from CAPM is below the median for the thick-market experiments is nowhere significant at the 10% level. That is, there is no statistical distinction between the CVT and thick-market experiments, despite the fact that the latter were ran with up to 63 subjects, whereas the former used at most 15 subjects.

---

[14] If the median fell between two outcomes, the lower of the two outcomes is reported instead of the median. This will bias our tests towards finding no improvement in equilibration in the CVT-based experiments.

[15] In addition, the numbers reported in Table 3 for the "median" average distance from CAPM in the parallel-market experiments was below the 50th percentile if the latter did not coincide with a realization, while the statistical tests took them invariably to be the 50th percentile. This means that our tests are biased towards finding no difference between the CVT and thin, parallel-market experiments. Also, in one of the eight thin-market experiments, the aggregate risk was substantially less than in the others (and compared to the CVT experiments), so that Sharpe ratios were generally much lower.

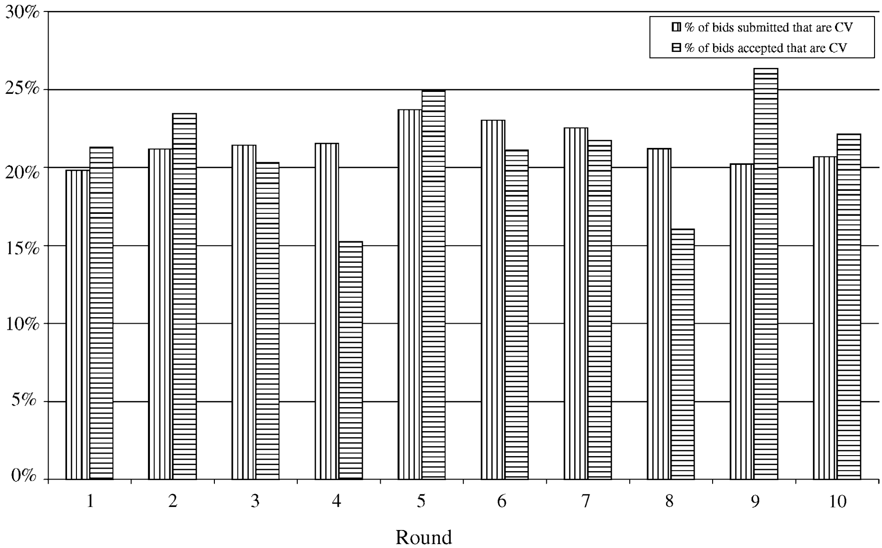

Fig. 3. Percent of bids submitted and accepted that involve at least two securities (denoted combined-value or CV), by round.

Overall, therefore, one can conclude that the CVT mechanism allows markets to move significantly closer to CAPM equilibrium pricing than the traditional, parallel-market continuous double auction with similar number of subjects. It manages to get markets equally close to CAPM equilibrium pricing as substantially more liquid markets arranged as parallel, continuous double auctions.

4.4. Liquidity enhancement through combined-value orders

Our last piece of evidence on the performance of the CVT mechanism concerns order submission. How frequently did subjects exploit the unique features of the CVT trading mechanism? That is, how often did they submit combined-value orders, those that involved more than one security? The second question is whether such packaged orders enhanced liquidity. In particular, what fraction of combined-value orders ended up being executed as part of the optimal allocation computed by the CVT mechanism on the basis of all the orders in the book?

To answer the first question, Figs. 3 and 4 report the percentage of submitted orders that were combined-value (defined to be orders for at least two securities, as opposed to orders for one security against cash only). Both by round (Fig. 3) and by period (Fig. 4), between 20% and 30% of the orders are combined-value. Consequently, subjects did exploit the flexibility that the CVT trading mechanism offered beyond the traditional system of parallel markets.

The success of our portfolio trading mechanism is further gauged in the finding that combined-value bids are no less likely to transact than are single-asset bids. Figs. 3

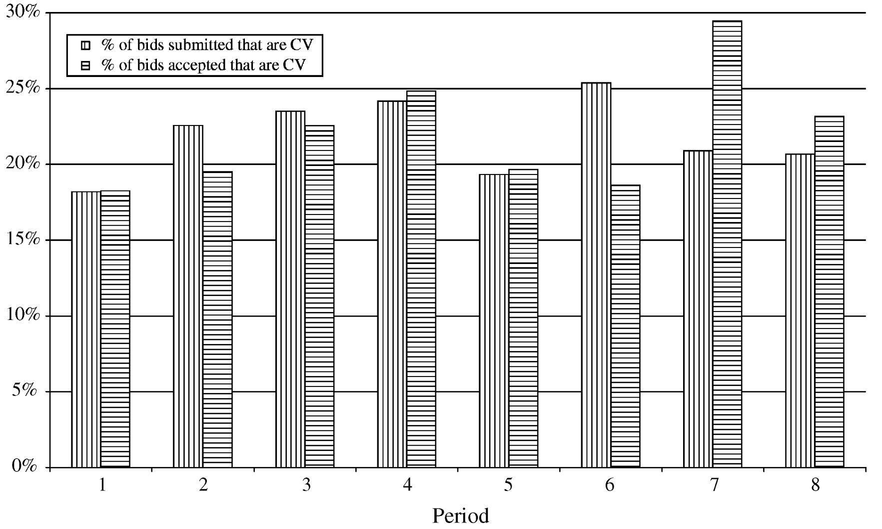

Fig. 4. Percent of bids submitted and accepted that involve at least two securities (denoted combined-value or CV), by period.

and 4 report that the percentage of orders accepted and executed by the CVT mechanism that are combined-value is about equal to the percentage of orders submitted as combined-value. Hence, combined-value orders not only enlarge the book, they also provide liquidity.

These results indicate that subjects seem to understand the CVT system. Not only did they submit orders to trade portfolios, but they did so that a large fraction became executable.

5. Conclusion

We reported results from seven small-scale experimental financial markets where order submission and trading took place through a portfolio trading mechanism – the combined-value trading (CVT) system. The results were compared to those from earlier experiments when markets were organized as a set of parallel double auctions. After a few rounds, the number of executed orders dropped, suggesting that gains from trade became exhausted and markets equilibrated. More pointedly, the new mechanism brings markets significantly closer to CAPM equilibrium pricing than the usual design of parallel, unconnected markets ($p < 1\%$). In fact, the average distance from CAPM equilibrium pricing is indistinguishable with that of large-scale experiments. Our findings cannot be attributed to chance: We reject that the CAPM provides no information to predict price changes, i.e., that prices move randomly and CAPM pricing obtains accidentally ($p < 1\%$). Finally, about 20–30% of the submitted *and* executed orders are combined-value (involve packages of at least two securities), indicating that subjects exploited the unique flexibility of the CVT trading mechanism, and that this generated liquidity.

The link between combined-value trading, liquidity and equilibration is suggested in asset pricing theory, which posits that agents are not interested in securities individually, but in portfolios (packages of securities). Unconnected, parallel double auctions do not allow agents to readily trade up to desired portfolio compositions, unless markets are sufficiently thick. The execution uncertainty of order submission in thin, parallel markets may even cause subjects to submit fewer orders, thereby further reducing liquidity, and keeping markets farther away from full equilibration. A comparison between our CVT experiments and others based on the same design but using the standard unconnected-markets trading mechanism corroborates these theoretical conjectures.

As an interesting by-product, our findings have implications for asset pricing theory, where illiquid assets are often thought of as generating an "equilibrium liquidity premium" over and above the usual risk premium. In view of the results of this article, it is odd to think about an equilibrium liquidity premium, because illiquid markets appear to be associated with markets that are kept from full equilibration. Because of this, it is hard to envisage an "equilibrium" liquidity premium.

Acknowledgements

Financial support was provided by The California Institute of Technology, The New Millennium Project of the Jet Propulsion Laboratory, and by the IBM Fellowship in Experimental Economics. The article benefited from comments during the May 2000 conference on liquidity at UCLA and during the September 2000 CEPR conference on finance in Manresa (Spain), as well as from suggestions from the referees and the editors, and from Chester Spatt.

Appendix A

Nonexistence of a price vector that satisfies (2)–(4) where quantities are determined as a solution to (1) with the additional constraint (5) can occur when subjects submit bids with flexibility levels (F_i) other than zero, because it generates nonconvexities in a player's revealed preferences. As an example, consider the following, Table 4, which details a bid schedule for a single asset.

Bid 7 is to sell 3 units for at least 3 francs per unit and is an inflexible order ($F_7 = 1$). Bid 2 is to pay up to 4 francs per unit for 3 units and is fully flexible ($F_2 = 0$). Bid 3 is to pay up to 2 francs per unit for 2 units, also flexibly. Surplus is maximized, given the flexibility constraints, if all bids except 4 and 8 are filled (3 is partially filled). There is, however, no competitive equilibrium, i.e., no solution to (6). To see this, notice that at any price above 2 francs bidder 3 is unwilling to buy units, but at any price below 3 per unit bidder 7 is unwilling to sell any units. There is no price so that demand equals supply, because the supply function effectively jumps where the demand function would cross.

Table 4

Bid number (i)	Units (q_i)	Price (b_i)	F_i[a]
1	2	5.0	0
2	3	4.0	0
3	2	2.0	0
4	2	1.0	0
5	−1	−0.5	0
6	−2	−1.5	0
7	−3	−3.0	1
8	−3	−4	0

[a]Flexibility, equals 0 if full, and equals 1 if none.

To determine prices and allocations when a competitive equilibrium does not exist, we first calculate a fully flexible equilibrium allocation by solving (1) and determine corresponding equilibrium prices. In the example in the above table, the fully flexible equilibrium price would be anywhere between 3 and at most 4 francs, say 3.5. The equilibrium quantity would be 5 units.

But bidder 7 submitted an inflexible order. At 3.5 francs, he is willing to supply 3 units but no less. Which means that the total supply at 3.5 francs is in fact 6 units, one more than in the flexible competitive equilibrium. The 6th unit will have to be sold to the 3rd bidder. Bidder 3 is, however, willing to pay no more than 2 francs.

Still, we can accommodate the inflexibility of bidder 7 by lowering the price that he gets and using the surplus to subsidize bidder 3. Bidder 7 is given only 3 francs per unit. He supplies 3 units, and the 3rd unit is allocated to bidder 3, who is charged 2 francs. The inframarginal bidders (numbers 1, 2, 5 and 6) each pay or receive the flexible equilibrium price, namely, 3.5 francs. Because of this, 3.5 francs is paid for 5 units (the total demand of bidders 1 and 2). At the supply side, 3 units are provided at 3.5 francs per unit. The other 2 units come from bidder 7, who is paid only 3 francs per unit, leaving an excess of 1 franc in total (2×0.5). This franc is used to subsidize trade for an additional unit. Bidder 3 is allocated the additional unit, and is asked to pay 2 dollars. Bidder 7 provides the additional unit, and receives 3 dollars for it.

Computation of optimal allocations that accommodate inflexibilities is a complex combinatorial problem. As in the example, we start from the flexible equilibrium and then attempt to determine a new allocation that satisfies binding inflexibility constraints. Subjects who cause binding inflexibility constraints are penalized as in the example above, and the penalty is used to subsidize additional trade.

Appendix B.

B.1. Description of thick, parallel-markets experiments

In virtually all respects, the thick, parallel-markets experiments against which the CVT experiments are compared (see Table 3) were identical. They were set up as

Table 5

Date	Subject number[a]	Signup reward (franc)	Endowments			Cash (franc)	Loan (franc)	Exchange rate ($/franc)
			A	B	Notes			
98/10/07	30	0	4	4	0	400	1900	0.03
98/11/16	23	0	5	4	0	400	2000	0.03
	21	0	2	7	0	400	2000	0.03
99/02/11	8	0	5	4	0	400	2000	0.03
	11	0	2	7	0	400	2000	0.03
99/04/07	22	175	9	1	0	400	2500	0.03
	22	175	1	9	0	400	2400	0.04
99/11/10	33	175	5	4	0	400	2200	0.04
	30	175	2	8	0	400	2310	0.04
99/11/11	22	175	5	4	0	400	2200	0.04
	23	175	2	8	0	400	2310	0.04

repetitions ("periods") of the same situation, with the same number of securities, states, and the same payoff matrix, namely:

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

Details of the thick-markets experiments can be found in the table below, to be contrasted with Table 1 (which provides the same information for the CVT experiments) (Table 5).

The four main differences between the thick-markets experiments and the CVT experiments were the following:

1. *The number of subjects*: As can be verified when comparing the above table with Table 1, the number of subjects in the thick-markets experiments was always higher (often a multiple) than in the CVT experiments.
2. *The trading interface*: In the thick-markets experiments, an electronic continuous open-book system was used, with one book per market, and where orders could not be made contingent on events in other markets. In contrast, the CVT experiments were based on an intermittent call market with a single electronic open book in arbitrarily specified portfolios.
3. *Randomization of initial allocations*: In the CVT experiments, each type of initial allocation was randomly assigned across subjects. In the thick-markets experiments, subjects were endowed with the same allocation across periods. Randomization of initial endowments tends to make price discovery more difficult, because subjects have to re-adjust each period.

Table 6

Date	Subject number[a]	Signup reward (franc)	Endowments A	Endowments B	Endowments Notes	Cash (franc)	Loan (franc)	Exchange rate ($/franc)
97/01/18	12	0	4	4	0	400	1500	0.03
97/01/27	10	0	4	4	0	400	1500	0.02
97/04/20	8	0	4	4	0	400	1500	0.02
97/04/30	5	0	4	4	0	400	1500	0.005
97/05/13	13	0	4	4	0	400	1500	0.03
97/05/19	11	0	4	4	0	400	1500	0.03
97/06/09	4	0	8	0	0	400	1500	0.03
	4	0	0	8	0	400	1500	0.03
	1	0	4	4	0	400	1500	0.03

4. *Location*: The thick-markets experiments were held over the internet, which means that the experimentors had no direct contact with the subjects. In constrast, subjects in the CVT experiments were all physically present in the SSEL laboratory at Caltech.

Further information and analysis of the thick, parallel-market experiments can be found in Bossaerts et al. (2001, 2002).

B.2. Description of thin, parallel-markets experiments

In virtually all respects, the thin, parallel-markets experiments against which the CVT experiments are compared (see Table 3) were identical. They were set up as repetitions ("periods") of the same situation, with the same number of securities, states. The payoff matrix was slightly different, namely:

Security	State X	State Y	State Z
A	75	475	50
B	200	125	275
C	200	200	200

In one experiment (97/06/09), the skewness of the payoff on security A was reversed:

Security	State X	State Y	State Z
A	275	275	50
B	200	125	275
C	200	200	200

Details of the thin-markets experiments can be found in Table 6, to be contrasted with Table 1 (which provides the same information for the CVT experiments).

The two main differences between the thin-markets experiments and the CVT experiments were the following:

1. *The trading interface*: In the thin-markets experiments, an electronic continuous double-auction system was used, with one auction per market, and where orders could not be made contingent on events in other markets. In contrast, the CVT experiments were based on an intermittent call market with a single electronic open book in arbitrarily specified portfolios.
2. *Randomization of initial allocations*: In the CVT experiments, each type of initial allocation was randomly assigned across subjects. In the thin-markets experiments, subjects were endowed with the same allocation across periods. Randomization of initial endowments tends to make price discovery more difficult, because subjects have to re-adjust each period.

Further information and analysis of the thin, parallel-market experiments can be found in Bossaerts and Plott (2002).

References

Athanasoulis, S., Shiller, R., 2000. The significance of the market portfolio. Review of Financial Studies 13, 301–329.

Bossaerts, P., 1999. The NYSE opening mechanism and portfolio trading. Caltech working paper.

Bossaerts, P., Plott, C., 2001. Basic principles of asset pricing theory: Evidence from large-scale experimental financial markets. Caltech working paper.

Bossaerts, P., Plott, C., 2002. The CAPM in thin experimental financial markets. Journal of Economic Dynamics and Control 26, 1093–1112.

Bossaerts, P., Plott, C., Zame, W., 2001. Prices and portfolio choices in financial markets, theory and experimental evidence. Caltech working paper.

Bossaerts, P., Plott, C., Zame, W., 2002. Prices and portfolio choices in financial markets: econometric evidence. Caltech working paper.

Bykowsky, M., Cull, R., Ledyard, J., 2000. Mutually destructive bidding: The FCC auction design problem Journal of Regulatory Economics 17, 205–228.

Cespa, G., 2001. A comparison of stock market mechanisms. Working paper, Università Pompeu Fabra.

Gallant, R.A., Tauchen, G., 1996. Which moments to match? Econometric Theory 12, 657–681.

Gouriéroux, A., Monfort, ., Renault, E., 1993. Indirect inference. Journal of Applied Econometrics 8, S85–S118.

Polk, C.W., Schulman, E., 2000. Enhancing the liquidity of bond trading. In: Rosen, J., Glisker, R. (Eds.), The Handbook of Fixed Income Technology. The Summit Group Press, New York.

Wohl, A., 1997. The feasibility of an index-contingent trading mechanism. Management Science 43, 112–121.

Wohl, A., Kandel, S., 1997. Implications of an index-contingent trading mechanism. Journal of Business 70, 471–488.