Price Formation in Multiple Simultaneous Continuous Double Auctions, with Implications for Asset Pricing

Elena Asparouhova*
Eccles School of Business, Uiversity of Utah

Peter Bossaerts†
Faculty of Business and Economics, University of Melbourne
Faculty of Economics, University of Cambridge

John Ledyard
The Humanities and Social Sciences, California Institute of Technology

October 12, 2021

*We would like to thank Bernard Cornet for pointing out a mistake in an earlier draft, and to Sean Crockett, Dan Friedman and Sophie Moinas for their comments and suggestions. We gratefully acknowledge comments on prior versions from participants in seminars (Bocconi, UBC, Copenhagen Business School, Columbia University, GATE (Lyon, France), U of Geneva, Hebrew University, U of Lausanne, U of Michigan, NYSE, New University in Lisbon, Norwegian Business School, Norwegian School of Economics and Business Administration, Ohio State University, SEC, Stanford, SIFR, Tel Aviv University, UC Berkeley, UC Irvine, UC San Diego, UC Santa Cruz, U of Kansas, UNSW, U of Paris–Dauphine; U of Vienna, U of Zurich) and conferences (2003 ESA meetings, 2003 WFAs, 2003 and 2009 SAET, 2004 Kyoto Conference on Experiments in Economic Sciences, 2005 Princeton Conference on Econometrics and Experimental Economics, 2005 Purdue Conference in honor of Roko Aliprantis, 2006 Decentralization Conference in Paris, 2010 Experimental Society Conference in Gothenburg, 2016 Experimental Finance Society Conference, Tucson), and the 2018 ESAM meetings in Auckland, New Zealand. We acknowledge the support from NSF grant SES-0527491 (Bossaerts, Ledyard), SES-0616431 (Bossaerts), SES-106184 (Asparouhova, Bossaerts), and the Swiss Finance Institute (Bossaerts).

†Corresponding author; Email: peter.bossaerts@unimelb.edu.au; Phone +33649490422.

Price Formation in Multiple Simultaneous Continuous Double Auctions, with Implications for Asset Pricing

Abstract

We propose a Marshallian model for price and allocation adjustments in parallel continuous double auctions. Agents quote prices that they expect will maximize local utility improvements. The process generates optimal allocations in the limit. In experiments designed to induce CAPM equilibrium, price and allocation dynamics are in line with the model's predictions. We identify, theoretically and empirically, a portfolio that is closer to mean-variance optimal throughout equilibration. This portfolio serves as a benchmark for asset returns even if markets are not in equilibrium, unlike the market portfolio, which only works at equilibrium. The theory has implications for momentum and liquidity.

Keywords: Continuous Double Auction; Walrasian Equilibrium; Marshallian Equilibration; Experimental Economics; Asset Pricing

1 Introduction

General equilibrium has become the widely accepted theoretical model for competitive markets and the benchmark against which those markets are empirically evaluated. A compelling reason to be interested in equilibrium is the "argument, familiar from Marshall, ... that there are forces at work in any actual economy that tend to drive an economy toward an equilibrium if it is not in equilibrium already."[1]

While there is wide consensus as to the appropriate equilibrium model, there is little consensus as to the "forces at work." Many models have been proposed, but none have been accepted as the appropriate canonical model. How the equilibrium prices and allocations are attained, and how, if at all, trading occurs out of equilibrium, remains to be discovered. The

[1] Arrow and Hurwicz (1958), p. 263.

lack of a consensus model of the forces that drive an economy towards equilibrium is a problem for applied economics, including policy analyses. If an inappropriate model is used in the design of economic policy, outcomes will not be as intended.

Until recently, attempts to settle this question have been mostly theoretical in nature[2] with no real evidence or philosophical foundation available to help sort the sensible from the inane. Traditional empirical analyses of markets shed no light on the processes because they do not have access to the fundamentals. But, with the advent and development of experimental economics, it is now possible to explore the forces that drive equilibrium.

The market organization we focus on in this paper is the continuous double auction (CDA) where individuals can submit bids (to buy) or asks (to sell) at any price, and whenever the highest bid is at a price at or above the lowest ask, a trade takes place immediately. In modern instances of the double auction, called the open-book system, bids and asks that are surpassed by more competitive orders (bids at a higher price or asks at a lower price) remain available, unless cancelled. The open book system is the preferred exchange mechanism of financial markets around the world, and in particular, of stock exchanges (NYSE, Euronext, LSE, NASDAQ, etc.). Recent advancements have been proposed where instead of immediate execution, there is a small interval over which orders accumulate in the book, called Frequent Batch Auctions (Budish, Cramton, and Shim, 2015). The model we propose also applies to those mechanisms.

It is well known from the experimental analyses of CDA markets (summarized in Crockett, 2013) that, in the first period of these experiments, (1) competitive equilibrium is not reached immediately – there is a process of adjustment – and (2) prices follow neither the

[2] Exceptions that study multiple simultaneous markets include the works of Plott (2001), Anderson, e.a. (2004) and Gillen e.a. (2021). These works report (price) dynamics that are in line with those reported here, as discussed later.

Walrasian tatonnement (whereby prices react to aggregate excess demands, but allocations are not adjusted) nor any of the various extant non-tatonnement theories (where allocations can also change). If the fundamentals and markets are repeated for additional periods, then (3) prices and allocations converge to their general equilibrium values and (4) between-period price changes follow the Walrasian tatonnement.

In this paper, we present a theory that explains the paths of prices and allocations within the first few periods of market experiments, before beliefs of likely paths could reasonably have been formed, and hence, where bets on their nature are pure speculation. It deserves emphasis that we model the paths *of allocations* as well as the paths of prices. The extant literature tends to focus only on price dynamics (Crockett, 2013).

There are three main assumptions underlying the theory. First, in the spirit of Marshall (1890), quantity moves to those offering the highest surplus to the market. Second, individuals quote prices that maximize their local utility gains taking the rules of engagement as given. Third, agents do not speculate, which means that they do not perceive drift in terms of trade that could improve their eventual allocations by postponing or accelerating transactions. Under these assumptions, the resulting offers are a convex combination of agents' marginal valuations and the prices.

The analysis is *not* on each bilateral trade separately as traditional CDA would require. Instead it invokes local market clearing,[3] defined as the transaction prices that cause net trades to sum to zero. In this sense, our model is more appropriate for the recently suggested frequent batch market mechanism (Budish, Cramton, and Shim, 2015).

Our theory is related to that of Friedman (1979), which itself follows up on the work of Smale (1976). Friedman identifies a process where allocations move in a Marshallian fashion:

[3] The local clearing prices are equal to the average of all offers.

throughout, prices are a weighted average of individuals' willingness-to-pay. Friedman (1979) focuses on stability and shows that the process converges to a Pareto-optimal allocation. However, the model misses detail on how offers are generated and how offers lead to trade. That is what our theory delivers.

Our theory is also related to that in Smith (1965) (see also Inoua and Smith, 2020). Smith shows that bids of many agents have an impact on prices and trades, not just those of the marginal agents, as in neoclassical accounts of Marshallian price adjustment (Samuelson, 1947). Our theory shares this prediction. In contrast to Smith's analysis, however, bids in our theory do not derive from Walrasian demand (or supply) functions. Instead, they result from agents' attempts to maximize local utility gains from trade.

To show the theory's power, we apply it to asset markets. It has a particularly intuitive appeal in the case of quasi-linear utility functions like mean-variance utility functions. Quasi-linear preferences naturally apply to the finance application of general equilibrium: the Capital Asset Pricing Model (CAPM) and its multi-factor generalizations (Roll, 1977).

We confront the finance application with data from nine experimental sessions, each with 6 to 8 replications ("periods") with varying parametrizations. The results provide strong support for the predictions regarding price and allocation dynamics. We test whether traditional Walrasian aggregate excess demands explain the remainder. We find that they do not. That is, Walrasian adjustment theory predicts neither price nor allocation dynamics.

The theory has important implications for empirical asset pricing, where for decades the concern has been to identify one mean-variance efficient portfolio, or a number of "factor portfolios" that add up to this efficient portfolio.[4] We find that price dynamics push one particular

[4]See Fama and French (2004). Since the set of mean-variance optimal portfolios is spanned by two portfolios, one of which necessarily is the risk-free security, it suffices to identify one additional mean-variance optimal portfolio to describe the entire set. See Roll (1977).

portfolio towards mean-variance efficiency throughout equilibration. Unlike in CAPM (equilibrium), it is not the market portfolio, but a risk-aversion weighted endowment portfolio. We refer to it as the Risk-Aversion Scaled Endowment Portfolio (RASE). In the experiments, we demonstrate that the RASE portfolio generates significantly higher average reward-to-risk ratios (Sharpe ratios) than the market portfolio.

The rest of the paper is organized as follows. The model setup and the theoretical results are presented in Section 2. Experimental methods are discussed in Section 3. Results are reported in Section 4. Implications for empirical asset pricing are in 5. Section 6 concludes.

2 Two Models of Market Dynamics

2.1 Preliminaries

2.1.1 The Economic Exchange Environment

Our analysis is done within the context of the standard model of pure exchange. There are I consumers, indexed by $i = 1, \ldots, I$. There are $K = 1 + R$ commodities, where the last R commodities are indexed by $k = 1, \ldots, R$, and the first is indexed by 0. We reserve this first commodity as a special one, and will designate it as the numeraire when needed.

Each individual i owns initial endowments $\omega^i = (\omega^i_0, \ldots, \omega^i_R)$, $\omega^i_k > 0$ for all i and k. $x^i = (s^i, r^i_1 \ldots, r^i_R)$ is the allocation of i. s^i is i's quantity of the numeraire commodity. $X^i = \{(s^i, r^i) \in \Re^K \mid r^i \geq 0\}$ is the admissible consumption set for i.[5] Each i has a quasi-concave utility function, $u^i(x)$. We assume that $u^i \in C^2$ (that is, u^i has continuous second derivatives) although many of our results would hold under weaker conditions. We also assume

[5] There is no lower bound on the numeraire.

that $\{x|u^i(x) \geq u^i(\omega^i)\} \subset \text{Interior}(X^i)$ and $u_0^i = \frac{\partial u^i(x^i)}{\partial x_0^i} > 0, \forall x^i \in X^i, \forall i$.

2.1.2 Time and the Continuous Double Auction

In a CDA experiment, traders begin with an endowment of commodities, ω^i. They proceed to make bids and offers over time. Often these are retained in a public book unless the trader decides to withdraw their bid or offer. A bid or offer in the book can be accepted by anyone. If accepted, trade occurs at that price. This goes on until a stopping rule is implemented. Although the CDA operates in continuous time, the intuition behind the theory is easier to understand in discrete time. Time is divided into discrete intervals of length Δ. With slight abuse of notation, the interval t is $[t, t+\Delta)$. $x_t^i = (s_t^i, r_t^i)$ denotes i's holdings at the beginning of interval t. Trade takes place and the change in i's holdings during interval t is $\Delta x_t^i = (\Delta s_t^i, \Delta r_t^i) = (s_{t+\Delta}^i - s_t^i, r_{t+\Delta}^i - r_t^i)$. $p_t = (1, q_t) \in \Re_+^K$ is the vector of K prices at which trades take place in interval t.

2.2 The Walrasian Model

Here we describe the standard Walrasian model of market dynamics as well as the variants known as non-tatonnement processes. There is nothing new here. We include this only as a reminder to the reader.

Given a price vector $p \in \Re_+^K$, the individual excess demand function of i is $\overline{e}^i(p, \omega^i) = \arg\max_{d^i} u^i(\omega^i + d^i)$ subject to $p \cdot d^i = 0$ and $\omega^i + d^i \in X^i$. The aggregate excess demand of the economy is $\overline{E}(p, \omega) = \sum_i \overline{e}^i(p, \omega^i)$, where $\omega = (\omega^1, \omega^2, ..., \omega^I)$.

Definition 1. *A price p^* and an allocation $x^* = (x^{*1}, \ldots, x^{*I})$ constitute a competitive equilibrium at $\omega = (\omega^1, \ldots, \omega^I)$ if and only if*

1. *Given prices* p^*, *the allocations* x^{*i} *are optimal:* $x^{*i} = \overline{e}^i(p^*, \omega^i) + \omega^i, \forall i$, *and*

2. *Markets clear; that is,* $\overline{E}(p^*, \omega) = 0$.

By Walras' law, we can limit our attention to the excess demands of all but the numeraire commodity, denoted $e^i(p, \omega^i)$ and $E(p, \omega)$, respectively. Also, since the price of the numeraire is fixed at 1, individual and aggregate excess demands can be written as $e^i(q, \omega^i)$ and $E(q, \omega)$, respectively, where $p = (1, q)$.

In Walrasian adjustment models, the main force driving price changes is the *tatonnement.* Prices of goods in excess demand (supply) go up (down). Let B be an $R \times R$ diagonal matrix with positive diagonal elements. The Walrasian tatonnement is:

$$\frac{q_{t+\Delta} - q_t}{\Delta} = BE(q_t, \omega) \tag{1}$$

$$x_t^i = \begin{cases} \omega^i & \text{if } E(q_t, \omega) \neq 0 \\ e^i(q_t, \omega^i) + \omega^i & \text{if } E(q_t, \omega) = 0 \end{cases} \tag{2}$$

The tatonnement is really only a model of prices since trades do not occur until prices have converged to their equilibrium values. (2) is not what is going on in most continuous markets where trading occurs as prices are changing.[6] Recognizing that, researchers have proposed many alternatives under the heading of *Non-Tatonnement* (NT) processes.[7]

An NT process works as follows. At the beginning of each time interval, agents know their individual holdings, x_t^i. Trade takes place during the interval at prices p_t. The holdings at the

[6] The tatonnement might describe, for example, the "book building" process in a call market if orders can be withdrawn (Plott and Pogorelskiy, 2017).

[7] See e.g. Negishi (1962), Uzawa (1962), Hahn and Negishi (1962).

end of the interval are $x^i_{t+\Delta}$. A new price is computed based on the excess demands at the price p_t and the holdings x_t. The Walrasian non-tatonnement dynamics are:

$$\frac{q_{t+\Delta} - q_t}{\Delta} = BE(q_t, x_t) \tag{3}$$

$$\frac{x^i_{t+\Delta} - x^i_t}{\Delta} = g^i(q_t, x^i_t), \tag{4}$$

where g^i is a vector function, $\sum_i g^i(q_t, x^i_t) = 0$, that also satisfies the Lipshitz condition. Different NT models impose different additional assumptions on the functions g^i, see Negishi (1962). In the CDA, there is no Walrasian auctioneer to set prices. There, (3) is interpreted as a predictive theory of prices: it predicts the price changes at $t + \Delta$ based on prices and allocations at t.

2.2.1 A Problem

In most multi-market CDA experiments, competitive equilibrium does not occur instantaneously except, perhaps, with replication in later periods. In addition, neither tatonnement, nor non-tatonnement dynamics fit the data.[8] A better theory is needed.

2.3 ABL Market Dynamics

Here, we describe a model based on Marshall's intuition but with a consistent micro-foundation. The model rests on four key hypotheses. The first captures the Marshallian intuition that **quantity moves to those individuals offering higher surplus to the**

[8]See Asparouhova, Bossaerts and Plott (2003), Anderson, e.a. (2004), Asparouhova and Bossaerts (2009), Gillen e.a. (2021), and Crockett (2013).

market. Let $b_t^i = (b_{1,t}^i, ..., b_{R,t}^i)$ be i's bid during the interval t. $b_{k,t}^i$ is i's stated willingness to pay (accept) to buy (sell) a unit of k in terms of the numeraire commodity 0.

Hypothesis 1. *Marshallian Trading*

$$\Delta r_t^i \,(= r_{t+\Delta}^i - r_t^i) = A(b_t^i - q_t), \quad i = 1, ..., I \tag{5}$$

where A is an $R \times R$ positive diagonal matrix and $A_{kk} = \alpha_k$, $k = 1, ..., R$.

In some markets, aggressive bidding attracts larger volume than in others. In this sense, α_k is a **liquidity parameter.** It is assumed that it does not vary over time.

The next two hypotheses are almost always requirements of a CDA system.

Hypothesis 2. *Instant Settlement (Payment with numeraire occurs at each trade)*

$$\Delta s_t^i = -q_t \cdot \Delta r_t^i \quad i = 1, ..., I. \tag{6}$$

Hypothesis 3. *Feasible Trading (Whatever is bought, is sold)*

$$\sum_{i=1}^{I} \Delta r_t^i = 0. \tag{7}$$

The last hypothesis, Hypothesis 4, specifies how individual traders determine their bids in a continuous double auction. It captures the idea that agents only consider small trades and do not speculate. Faced with the fact that large orders will move prices unfavorably, intractable strategic uncertainty, and a lack of futures markets and rational expectations, **agents make only small (local) adjustments to their holdings**. This can be motivated using game theory, but it is also a *fact* in field markets.[9] Faced with uncertainty about where prices will go next, **agents do not speculate**. They take current prices as given.

To motivate Hypothesis 4, assume traders only consider small local adjustments that maximize their gain in local utility Δu_t^i. For very small Δ, $\Delta u_t^i \approx \nabla u^i(x_t^i) \cdot (\Delta s_t^i, \Delta r_t^i)$ where $\nabla u^i(x_t^i)$ is the gradient of u^i at x_t^i. Under Hypotheses 1 and 2, the change in i's utility that results from a bid b_t^i at time t will be:

$$\Delta u_t^i \approx u_0^i(x_t^i)(\rho^i(x_t^i) - q_t) \cdot \Delta r_t^i = u_0^i(x_t^i)(\rho^i(x_t^i) - q_t) \cdot A(b_t^i - q_t),$$

where $\rho_k^i(x^i)$ denotes the marginal rate of substitution between commodities 0 and k for $k = 1, ..., R$ if i's holdings are x^i.[10] ρ_k^i represents i's marginal willingness to pay (or be paid) for units of k in units of commodity 0. $\rho^i(x^i) = (\rho_1^i(x^i), ..., \rho_R^i(x^i))$ and $\nabla u^i(x_t^i) = u_0^i(x_t^i)(1, \rho^i(x_t^i))$.

To locally optimize, i wants to choose b_t^i so that the direction of change of $x_t^i = (s_t^i, r_t^i)$ is proportional to the gradient. This means they want $A(b_t^i - q_t) = c^i\Delta(\rho^i(x_t^i) - q_t)$, where the parameter c^i is a characteristic of i. It determines the step size and rate of trading. Larger c^i imply a greater urgency to trade. We call this i's **impatience parameter** and assume it does not change over time.

Remark 1. *This behavior is incentive compatible in the following sense. If both the quantity adjustment rule, Hypothesis 1, and the price setting rule, Hypothesis 3, are known and taken as given, and $\alpha_k = \alpha$, for $k = 1, ..., R$, then there are $(c^1, ..., c^I)$ such that the bids derived*

[9] Financial markets have become more competitive, and trade sizes have decreased dramatically. "Splitting orders" has become an important concern in algorithmic trading. See Avellaneda, Reed and Stoikov (2011). Further empirical evidence that trade takes place "in smalls" can be found in O'Hara, Yao and Ye (2014). In a market with continuous order submission and trading, the small-orders assumption can easily be justified theoretically; see Rostek and Weretka (2015).

[10] $\rho_k^i(x^i) = \frac{\partial u^i(x^i)/\partial x_k^i}{\partial u^i(x^i)/\partial x_0^i}$.

above are a local Nash equilibrium.[11]

The final intuition behind Hypothesis 4 concerns the timing of information and actions. When an agent computes their bid at the start of interval t, they do not know q_t. They only know the prices and allocations at the end of the $t - \Delta$ interval. Because Δ is assumed to be very small, it is likely that bids at t are based on the prices and allocations arrived at in the interval $t - \Delta$.

Hypothesis 4. *Local Optimization and Lagged Prices*

$$b_t^i = q_{t-\Delta} + c^i \Delta A^{-1}(\rho^i(x_t^i) - q_{t-\Delta}), \forall i, \forall t > 0.$$

For the curious, Section B.1 of the Appendix contains a discussion of the model and its implications when q_t is used in place of $q_{t-\Delta}$ in Hypothesis 4. That model implies that bids and prices are *simultaneously* determined in the time Δ. The model is not consistent with the data, as explained in Appendix B.2.

This leaves the initial price, q_0, to be specified. The initial price is likely context-dependent and can plausibly equal the vector of mean payoffs in an asset pricing setup, be related to prices in the previous period when applied to replications of the same situation, or be equal to the average of the values of the initial endowments.

Hypothesis 5. *The initial price q_0 is some arbitrary positive vector.*

In our empirical analysis, the focus will be on price *changes*, so Hypothesis 5 is never in play.

Hypotheses 1-5 are the ABL model.

[11] This is similar to a result of Roberts (1979). A proof is provided in section A.1 of the Appendix.

Remark 2. *We have assumed that agents do not speculate. The beginning of an analysis under speculation can be found in Appendix C. Speculation becomes an important concern in later replications in an experiment, when these replications are identical, meaning participants have the opportunity to form beliefs about likely price dynamics. Here, we focus on early replications, or replications with varying parametrizations.*

The dynamics of the ABL model are straightforward. Entering interval t, consumer i has an allocation $x_t^i = (s_t^i, r_t^i)$ and knows the price from the previous interval $q_{t-\Delta}$. In the interval, bids are formed based on Hypothesis 4 and trade occurs at new prices based on Hypotheses 1-3. Prices adjust rapidly to ensure that trading, according to Hypothesis 1 and 2, adds up to zero (Hypothesis 3). Leaving the interval, trader i now owns $x_{t+\Delta}^i = (s_{t+\Delta}^i, r_{t+\Delta}^i)$ and knows the prices q_t. This process, given the initial price q_0, is formalized in equations (8)-(10).[12]

$$r_{t+\Delta}^i = r_t^i + \Delta\left(-\bar{c}(\bar{\rho}_t - q_{t-\Delta}) + c^i(\rho_t^i - q_{t-\Delta})\right) \tag{8}$$

$$s_{t+\Delta}^i = s_t - q_t \cdot (r_{t+\Delta}^i - r_t^i) \tag{9}$$

$$q_t = q_{t-\Delta} + \bar{c}\Delta A^{-1}(\bar{\rho}_t - q_{t-\Delta}) \tag{10}$$

where $\bar{c} = \frac{\sum_i c^i}{I}$ and $\bar{\rho}(x_t) = \frac{\sum_i c^i \rho^i(x_t^i)}{\sum_i c^i}$.

The limiting behavior of the dynamics is most easily seen in continuous time.[13] Dividing

[12] See Appendix A.2 for details.

[13] Convergence in continuous time implies that if step sizes, c^i, are not too large, then there will also be convergence in discrete time.

(8) and (10) by Δ and letting $\Delta \to 0$, we get the continuous time version, for $t > 0$:[14]

$$\frac{dr_t^i}{dt} = c^i(\rho_t^i - q_t) - \bar{c}(\bar{\rho}_t - q_t), \forall t > 0 \tag{11}$$

$$\frac{ds_t^i}{dt} = -q_t \cdot \left((c^i(\rho_t^i - q_t) - \bar{c}(\bar{\rho}_t - q_t)\right), \forall t > 0 \tag{12}$$

$$\frac{dq_t}{dt} = -\bar{c}A^{-1}(q_t - \bar{\rho}_t), \forall t > 0 \tag{13}$$

Remark 3. *When taking limits, one important subtlety of the ABL model is lost. The discrete-time equations specify dynamics over* two *intervals:* $[t-\Delta, t)$ *and* $[t, t+\Delta)$*. In continuous-time, everything collapses effectively to one interval. E.g., in discrete time, price changes over* $[t - \Delta, t)$ *depend on marginal rates of substitution at the* end *of the interval (i.e., at* t*); see (10). In continuous time, it does not matter whether marginal rates of substitution are based on holdings at the beginning or end of an interval, because adjustment is smooth. To preserve discrete-time subtleties, one could add random shocks to the adjustment, and appeal to Itô calculus. Limit (Itô) processes are* not *smooth (time series are nowhere differentiable with respect to time). Consequently, timing subtleties from discrete time are retained in continuous time. As reported in Section 4 below, the discrete-time subtleties matter empirically. Price changes within observation intervals in our trading sessions are driven by holdings at the* end *of each such interval, as predicted by the ABL model. The Walrasian model, in contrast, predicts that price changes are based on (excess demands computed from)* lagged *holdings. The Walrasian model fails if only because of timing issues. Timing is an under-appreciated dimension in which Marshallian and Walrasian dynamics differ. In Marshallian dynamics, prices are determined by current willingness to pay; in Walrasian dynamics, prices are determined by past excess demands. This subtle but important difference in the models will be crucial for our empirical*

[14]See Appendix A.3 for details.

work.

There is an analogy to the First Welfare Theorem of General Equilibrium Theory: the allocation at any rest point is a Pareto-optimal allocation. By the Second Welfare Theorem the rest point is also a competitive equilibrium at that allocation. If there are no income effects, the continuous process (11)-(13) will converge to a rest point from any initial price and allocation. This may not be true for the discrete process (8)-(10) if step sizes are too large.

Theorem 1. *Convergence to Pareto Optimal Allocations*[15]

If (i) there are no income effects, i.e., $u_0^i(x^i) = 1$ *for all* i *and all* $x^i \in X$*, and (ii)* $r_t^i > 0$ *for all* t*, then for the dynamics in (11) - (13),* $(x_t, p_t) \to (x^*, p^*)$ *where* x^* *is Pareto-optimal and* (p^*, x^*) *is a competitive equilibrium at* x^**.*

Remark 4. *Along the path generated by (11) - (13), it is possible that* $du_t^i/dt < 0$*. With the bidding lag,* $du_t^i/dt = u_{0,t}^i((\rho^i(x_t^i) - q_t) \cdot c^i(\rho^i(x_t^i) - q_t) - \sum_k(\rho_k^i(x_t^i) - q_{k,t})\alpha_k(dq_{k,t}/dt))$ *While the first term is non-negative, the second term is not necessarily so. Traders basing their bids on lagged prices do not anticipate and cannot protect themselves from "ex post" adverse trades. For example, if prices are rising fast, slow agents may trade into increasing prices when they want to buy.*

Remark 5. *The possibility that* $du_t^i/dt < 0$ *(among other differences) distinguishes the ABL theory from Friedman (1979) and Smale (1976). Specifically, our allocation dynamics do not satisfy Friedman's condition (P).*

[15] The proof of this theorem is relegated to Section A.4 of the Appendix.

2.4 Comparing Walrasian vs. ABL Dynamics

The Walrasian and ABL models can imply significantly different paths of price adjustment. This can be seen in the simple example in Figure 1. There $R = 1$ and $I = 2$, utility functions are quasi-linear (the inverse demand functions therefore equal the marginal rates of substitution ρ), and the aggregate endowment is $W = r^1_{t-\Delta} + r^2_{t-\Delta} = r^1_t + r^2_t$. We measure the holding of trader 2 from right to left starting at W. The competitive equilibrium allocation and the resting point of the ABL model occur where ρ^1 and ρ^2 cross, with q^e denoting the equilibrium price.

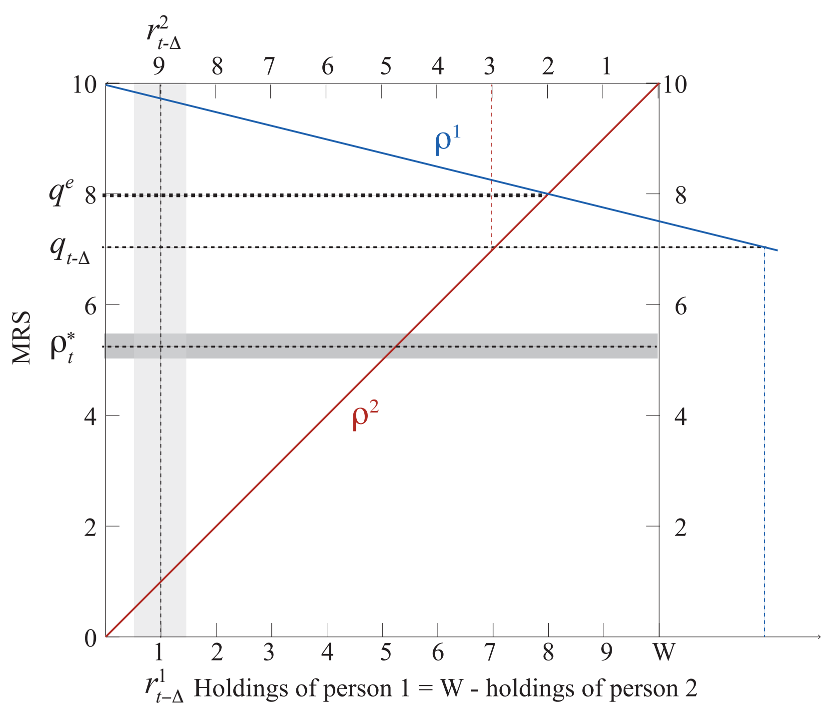

Figure 1: MRS (Marginal Rate of Substitution ρ^i) in a 2-commodity, 2-person economy, as a function of holdings of agent 1. Equilibrium price equals q^e. Last traded price equals $q_{t-\Delta}$. The Walrasian equilibration model predicts that the price will increase because, at $q_{t-\Delta}$, there is excess demand: agent 2 demands three units and agent 1 demands more than W units, while total supply equals only W units. In contrast, ABL predicts that the price will decrease, to ρ^*_t, which equals the average of the ρ^is at current holdings.

In Figure 1, $r^1_{t-\Delta}$ denotes 1's holdings at $(t - \Delta)$, while 2 holds $r^2_{t-\Delta} = W - r^1_{t-\Delta}$. The most recent price, $q_{t-\Delta}$, is below the equilibrium price. At the given holdings, and given the most recent price, there is excess demand for the good (at $q_{t-\Delta}$, individual 2 demands 3 units, and 1 demands more than W units) so the Walrasian model requires the price to *increase*, i.e., $q_t - q_{t-\Delta} > 0$. To determine the sign of $q_t - q_{t-\Delta}$, the ABL model uses the allocations at t, r_t^1 and r_t^2. Given small changes in quantities, these allocations will be close to $r_{t-\Delta}^1$ and $r_{t-\Delta}^2$, as depicted by the vertical band. As a result, the average weighted marginal rate of substitution, $\rho_t^* = \bar{\rho}(r_t)$, falling in the corresponding horizontal band, is lower than the price $q_{t-\Delta}$ meaning the ABL model predicts that the price would *fall*, i.e., $q_t - q_{t-\Delta} < 0$.

The difference in the implications of the two models when $R > 1$ is also very stark if we restrict attention to a very special environment: the Capital Asset Pricing Model (CAPM). The CAPM is theoretically simple and is of its own interest since it serves as the foundation of both asset market experiments and empirical analyses on historical data from the field. In the CAPM, all utility functions are of the form:

$$u^i(x^i) = s^i + \mu \cdot r^i - (a^i/2)(r^i) \cdot (\Omega r^i), \tag{14}$$

where μ is an R-dimensional vector of positive constants, Ω is an $R \times R$ positive-definite matrix of constants, and a^i is a positive scalar constant. In asset pricing models, μ is interpreted as the expected payoff of an asset, Ω is the payoff covariance matrix across the assets, and a^i is a measure of risk aversion. For these utility functions,

$$\rho^i(x^i) = \mu - a^i \Omega r^i \quad \text{and} \quad e^i(q, x^i) = \frac{1}{a^i}\Omega^{-1}(\mu - q) - r^i. \tag{15}$$

Combining (15) with (10) yields:[16]

$$\frac{q_t - q_{t-\Delta}}{\Delta} = A^{-1}\Omega \frac{\sum c^i a^i e^i(q_{t-\Delta}, x^i_t)}{I} \tag{16}$$

Comparing (16) with (3), we can see three fundamental differences between the *price dynamics* of the ABL model and those of the Walrasian model in the CAPM environment.[17]

1. **Cross-Security Effects Emerge.** In the ABL model, changes in the price of commodity k depend not only on the excess demand for k (as in the Walrasian model) but also on the excess demand of the other commodities. For example, if the off-diagonal entries of Ω are negative (indicating the commodities are complements),[18] the excess demand for $j \neq k$ puts upward pressure on the price of k. This means that the price of k could increase, even though there is an excess supply of it. This cannot happen under Walrasian price dynamics.

2. **Heterogeneity in Risk Aversion, Impatience and Liquidity Matters.** In the ABL model the excess demand functions of traders with higher $a^i c^i$ are weighted more heavily in how they affect the changes in prices. The desires of the more risk averse and the more impatient thus have a larger impact on price changes. In the Walrasian model it is the less risk averse who have a larger impact on price changes.

3. **Timing Is Different.** See Remark 3. In the Walrasian model, prices in interval t are determined by prices and allocations in period $t - \Delta$. In the ABL model, prices in period t are determined by prices in period $t - \Delta$ and by *allocations in period t.*

[16]See section D of the Appendix for the details of the derivation.

[17]The premultiplication by Ω of the excess demands might remind some of the Newton-Raphson algorithm. We discuss this in section E of the Appendix.

[18]A similar analysis applies when the commodities are substitutes or when there is a mix of both.

The three differences are testable in the lab and motivate the design of our experiment.

As to *allocation dynamics*, using (15), the following system of difference equations describes agent-level changes in allocations:

$$\frac{r^i_{t+\Delta} - r^i_t}{\Delta} = -\Omega\left(c^i a^i r^i_t - \frac{\sum c^i a^i r^i_t}{I}\right) + \left(c^i - \bar{c}\right)(\mu - q_{t-\Delta}) \tag{17}$$

In ABL, the changes in an agent's allocations depend on (i) how far impatience and risk-aversion scaled holdings are from the average impatience and risk-aversion scaled holdings, plus (ii) the differences between expected payoffs and lagged market prices, provided the agent's impatience is different from the average. The second term disappears if impatience is the same across agents; the first term remains under equal impatience, as long as risk aversion is heterogeneous. The covariance matrix pre-multiplies the first term. As a consequence, ABL predicts cross-security effects in allocation dynamics in the same way it predicts them in price dynamics. The effects are opposite for prices and allocations however, because of the negative sign in front of the first term of (17).

Equations (16) and (17) will form the basis of our empirical analysis.

3 Experimental Methods

3.1 Framework

Our experimental design builds on the CAPM. Agents have mean-variance preferences with fixed risk-to-reward trade-offs, and hence, no wealth effects. Prior experiments have shown robust convergence to equilibrium; see Asparouhova, Bossaerts and Plott (2003); Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007).

CAPM predicts that, in equilibrium, one particular portfolio is mean-variance optimal. This portfolio is the market portfolio. In CAPM, agents' *total* demands (holdings plus excess demands) are the same for all agents, up to a constant of proportionality equal to the inverse of risk aversion. This is obtained by rewriting (15):

$$r_t^i + e^i(q_t, x_t^i) = \frac{1}{a^i}\Omega^{-1}(\mu - q_t). \tag{18}$$

The property is known as "portfolio separation." As a result, in the Walrasian equilibrium, the right-hand-side must equal to the total supply of assets, i.e., the "market portfolio." The market portfolio is defined as the per-capita endowment portfolio of risky assets, with holdings equal to $\overline{r} = \frac{1}{I}\sum_{i=1}^{I} r^i$. Consequently this means that, in equilibrium, the market portfolio must be mean-variance optimal, for otherwise it would not be proportional to agents' demands. See Roll (1977).

Equilibrium prices are as follows.[19]

$$q^* = \mu - \frac{1}{\frac{1}{I}\sum_i \frac{1}{a^i}}\Omega\overline{r}. \tag{19}$$

In the laboratory, CAPM works well; see, e.g., Bossaerts and Plott (2004); Bossaerts, Plott and Zame (2007). Here is an example, from a classroom session in an advanced investments class at the University of Melbourne. Forty-eight students were asked to trade to maximize

[19] It is straightforward to check that, at these prices, the sum of the individual excess demands (18) equals zero, and hence, markets equilibrate. When converted to restrictions on returns (payoffs divided by prices), the equation becomes the well-known requirement that expected returns in excess of the risk-free rate be proportional to the covariance of returns with those on the market portfolio.

their payoffs given by (14), with $a^i = 0.01$, for all i,

$$\mu = \begin{bmatrix} 5 & 5 & 5 \end{bmatrix} \text{ and } \Omega = \begin{bmatrix} 16 & -5 & -14 \\ -5 & 16 & 9 \\ -14 & 9 & 16 \end{bmatrix}.$$

Notice that mean-variance preferences are *induced* by asking students to directly optimize the CAPM payoff function. In the sequel, we will nevertheless refer to μ as the vector of expected payoffs, and Ω as the covariance matrix.

The three securities had equal expected payoffs and equal variances. But in equilibrium prices differ because (i) supplies were unequal, with the third security being in the shortest supply and (ii) the first security had negatively correlated payoffs with the others, while the other two had positively correlated payoffs. Equilibrium prices were:

$$q^* = \begin{bmatrix} 5.125 & 1.5 & 3.5 \end{bmatrix}.$$

The equilibrium price of the third security is not the highest even if it is in the shortest supply. The intuition is simple: the first security, with the highest equilibrium price, is more valuable because its payoff is negatively correlated with that of the others. Participants were not told the per-capita supplies. Hence, even if they knew CAPM, they could not possibly compute equilibrium prices.[20]

Trade in this sample laboratory market took place in an online continuous open-book trading platform (called Flex-E-Markets[21]). Participants could submit limit orders for any

[20] The results of a quick poll before trading confirmed that most participants expected prices to be equal.
[21] See http://www.adhocmarkets.com.

of the securities for the duration of the class exercise (about 35 minutes). Participants were provided with a tool that evaluated the performance of their current portfolios as well as the net performance of any trades they wished to make.

Figure 2 displays the evolution of trade prices, during the first replication, of the three risky securities (referred to as Stock A, Stock B and Stock C). Prices convincingly evolved from expected values to equilibrium levels.[22]

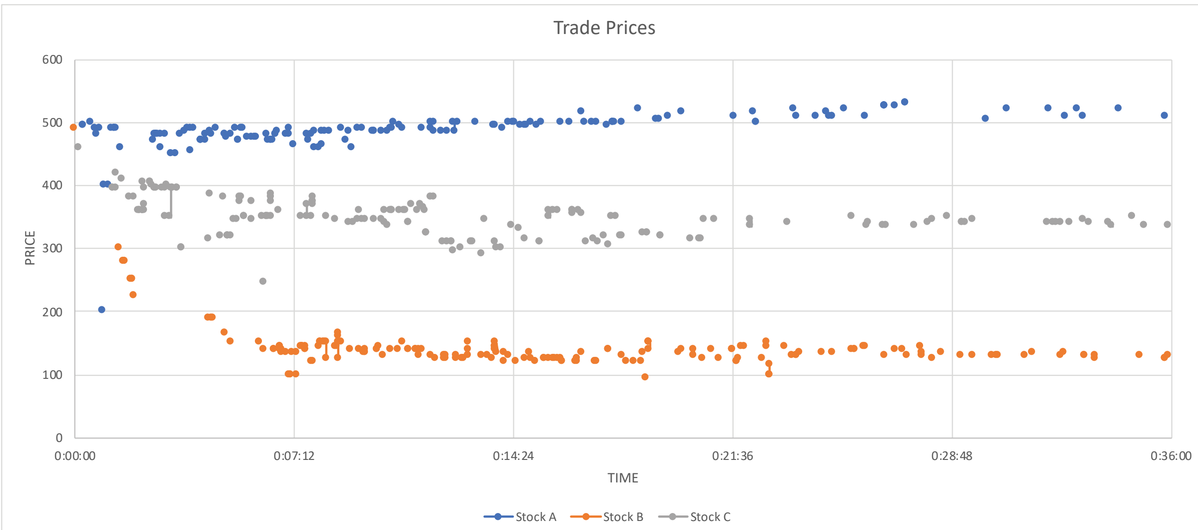

Figure 2: Transaction prices (in cents) during a class experiment. Forty-eight participants traded three risky securities ("Stocks" A, B and C) with known, equal payoff distributions but different, unknown total supplies. Predicted equilibrium prices, in cents: 513 (A; blue), 150 (B; orange) and 350 (C; grey).

Participants were divided into three groups based on their initial portfolio allocations. They only knew their own allocation. The first group started with 15 of the first security and none of the other securities; the second group started with allocations of 9, 20 and 0, and the third group started with 0, 10, and 18. In equilibrium, they should all end up with the same allocation, since they all faced the same risk aversion parameter. Final allocations necessarily equal the market portfolio. Figure 3 plots the evolution of the *difference* of the

[22]We were agnostic as to the price levels markets would start from; see Hypothesis 5. In the experiment, prices started from expected value. That is, $q_0 = \mu$.

per-capita holdings of Group 1 and the market portfolio, over intervals of 5 trades each. The

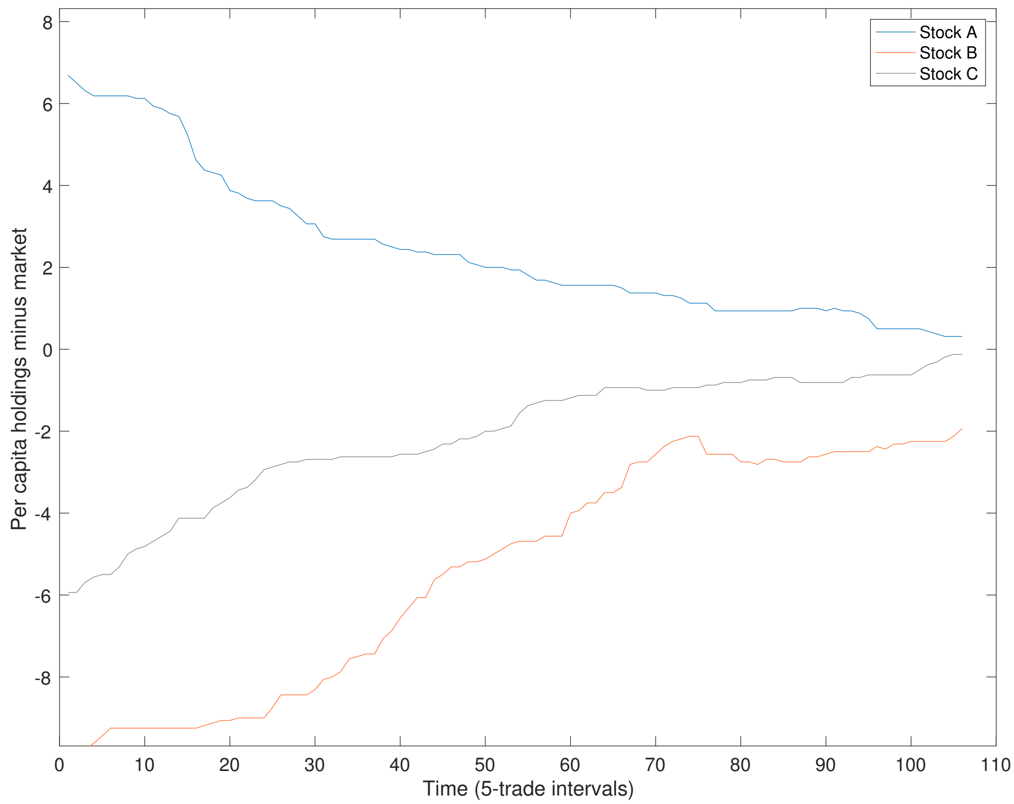

Figure 3: Evolution of differences between (i) per-capita holdings of A (blue), B (orange) and C (grey) of the first group of participants, and (ii) the market portfolio. Initial holdings are 15 units of A each and 0 of B and C. The market portfolio consisted of (per capita) 8 units of A, 10 of B and 6 of C. Differences converge to zero, implying that per-capita holdings converged to CAPM predictions. Time is measured in intervals of 5 transactions.

figure shows how per-capita holdings gradually move towards the equilibrium level. Notice that the evolution is far more gradual than the price evolution.

In the class experiment, we induced mean-variance preferences, by tying performance directly to the CAPM utility function in (14). There was no explicit uncertainty in the experiment; performance (payoffs) were immediate once allocations were known. We could also have drawn payoffs from distributions with mean μ and covariance matrix Ω, but then we would not have controlled the risk aversion parameter, so we could not have unambiguously derived equilibrium price levels. In addition, we would have to make the auxiliary assumption that mean-variance preferences explain choices in the experiment.[23]

As with the classroom experiment presented above, the experimental sessions we ran to test the theory of this paper also relied on induction of mean-variance preferences. To simplify the setup, the experiments had two, not three, risky securities.[24] Also, since the theory has predictions for economies with heterogeneous risk aversion, we varied the risk aversion coefficient across subjects.

3.2 Hypotheses

The theory makes precise predictions about the evolution of *prices* as well as *allocations*. Allocation changes depend on risk aversion and are therefore analyzed as average changes in holdings across subjects who belong to homogeneous groups. Groups are defined by initial allocations and/or risk aversion coefficients. The parameters μ and Ω in the payoff functions are the same regardless of group.

Prices

For the mean-variance utility functions in (14), individual marginal willingness to pay is $\rho^i(x^i) = \mu - a^i \Omega r^i$, while risk-aversion weighted average willingness to pay is $\bar{\rho}(x) = \mu - \Omega \frac{\sum_i a^i c^i r^i}{\sum_i c^i}$. Hence, the price dynamics implied by our model ABL, in discrete time, are given by the equation $q_t - q_{t-\Delta} = \bar{c} \Delta A^{-1} \left(\mu - \Omega \frac{\sum_i a^i c^i r^i_t}{\sum_i c^i} - q_{t-\Delta} \right)$. See Online Appendix D, Equation OA.18. In the sequel, we set $\Delta = 1$.

[23]When introducing uncertainty explicitly, Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007) show that mean-variance preferences provide only a crude approximation of individual choices, even if CAPM accurately predicts prices.

[24]Online Appendix F.2 reports results from earlier sessions with three risky securities, but where mean-variance preferences were not induced.

Since we would like to compare this to the Walrasian model (3), we want to write it in terms of excess demand functions, as in (16):[25]

$$q_{t+1} - q_t = \frac{1}{I} A^{-1} \Omega \left(\sum_i a^i c^i e^i(q_t, x^i_{t+1}) \right). \tag{20}$$

The equations summarize the price dynamics under ABL. They constitute the key hypothesis which we test on the data. They link changes in prices to Walrasian excess demands. As discussed in the theory section, there are three unusual aspects compared to the traditional Walrasian adjustment model. We repeat them here for convenience.

1. The covariance matrix Ω pre-multiplies the vector of risk-aversion weighted excess demands. This means that the *excess demand of one security determines price changes of all other securities*, and the effect is proportional to the corresponding payoff covariances.

2. Excess demands are *weighted* by risk aversion, liquidity and impatience parameters. In our experiments, the liquidity and impatience parameters will not be controlled, so we will assume that they are the same for everyone.[26]

3. Excess demands are evaluated at *end-of-period* holdings, unlike in the Walrasian model (3). We already emphasized this subtle difference in timing between the two models; see Remark 3.

In the empirical tests, we will pay close attention to these three features. To directly test the first feature, we pre-multiply the vector of risk-aversion weighted excess demands by the covariance matrix, so that cross-security effects are no longer present. That is, we run the

[25]Equation (16) specifies price changes over period $t - \Delta$ while Equation (20) does the same over period t.

[26]There is evidence that impatience relates to risk aversion, however: see Asparouhova and Bossaerts (2009). We will return to the issue in the Results section; see the discussion concerning Figure 7.

following multi-equation regression:

$$q_{t+1} - q_t = B \text{ WE}(q_t, \{x^i_{t+1}, \text{all } i\}) + \epsilon_t, \tag{21}$$

where $\text{WE}\left(q_t, \{x^i_{t+1}, \text{all } i\}\right) = \Omega\left(\frac{1}{I}\sum_i a^i e^i(q_t, x^i_{t+1})\right)$.

The main restriction is that the coefficient matrix B is *diagonal*. We cannot say much about the magnitude of the diagonal coefficients *except that they should be strictly positive*. In (21), an error term ϵ_t is added, to reflect noise in the dynamics. In the empirical analysis constant terms will also be added. These will be period-specific if the data straddle multiple periods.[27]

Let us illustrate the regressions in (21) using the class experiment. Figure 4 displays scatter plots of price changes and the regressors. Price changes were computed over intervals of five transactions. The vertical axes in the figure reproduce the price changes from Figure 2, over intervals of five trades. The horizontal axes display the regressors in (21), also calculated every five trades. The prediction is that there is a positive relationship between price changes of a security i (= A, B, C) only for regressors WE(i). That is, the relation is (strictly) positive only for the plots on the diagonal, where observations are plotted in red. No relationship should exist in plots off the diagonal, where observations are plotted in blue. Visual inspection suggests that this is indeed the case. A formal test of the hypothesis is provided above each of the plots. Displayed is the magnitude of the estimated slope coefficient, as well as the corresponding z-statistic. z-statistics beyond 2 can be considered "significant" ($p = 0.02$).

[27]The constant term plays no role in the theory, but may be needed empirically to avoid model misspecification. If our model does not explain everything in the data (as one should expect), imposing zero intercepts can lead to serious biases in the estimation of slope coefficients, and hence, mis-interpretation of the findings.

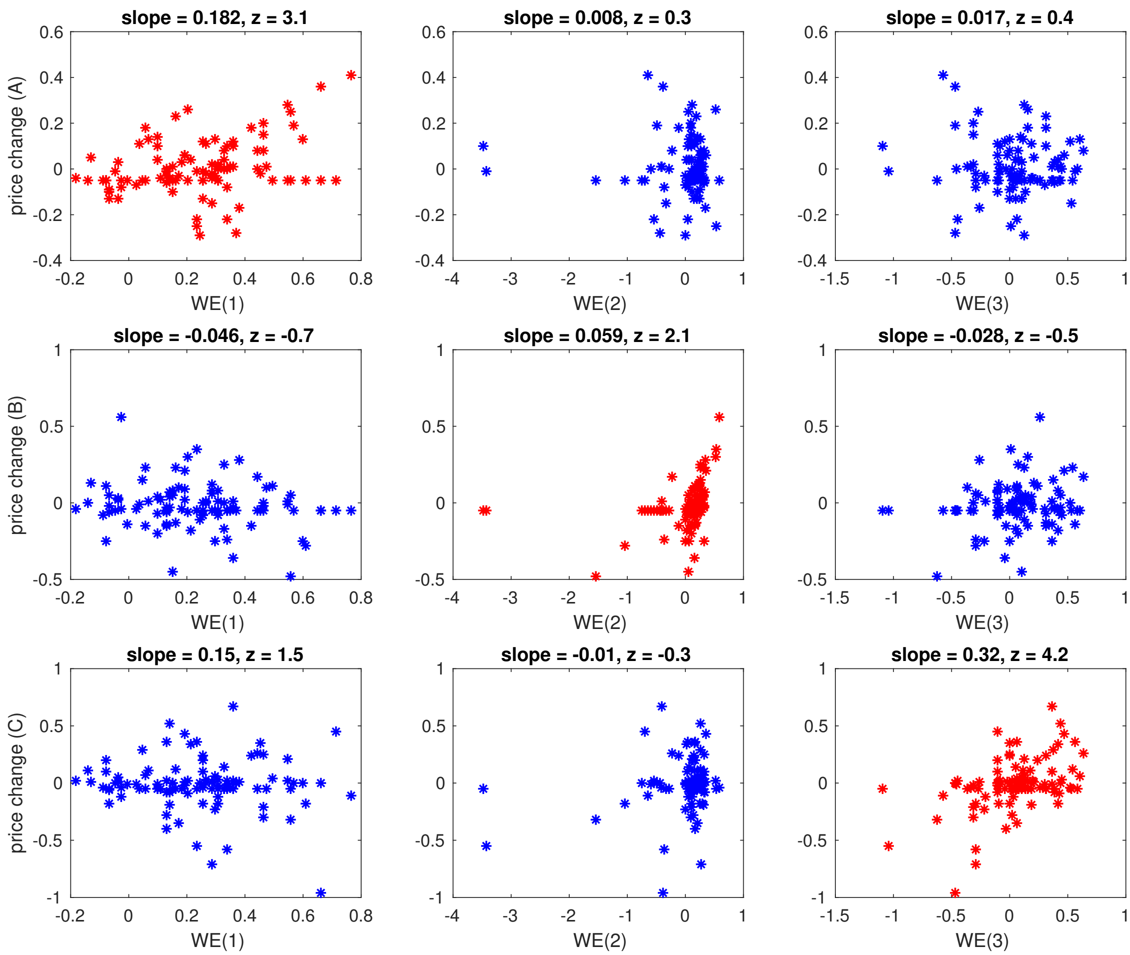

Figure 4: Scatter plots of price changes over intervals of five (5) trades against regressors in (21). Estimates of regression slopes (using Huber's robust regression) and corresponding z-statistics are indicated on top of each plot. The ABL model predicts that the plots with red observations should generate a strictly positive slope, while the remaining plots should have zero slopes. The results are consistent with the ABL model (using $p = 0.02$). Number of observations per plot: 101.

Slope coefficients are estimated using Huber's robust regression with $\delta = 2.0$.[28] Consistent with the theoretical predictions, slope coefficients on the diagonal are all significant, while none in the off-diagonal plots are.

Walrasian dynamics are different. From (3), the price-change regressions for the Walrasian

[28]Huber's robust regression uses a loss function that treats outliers differently compared to least squares. With parameter δ, the loss function is defined as: $L(\epsilon) = \epsilon^2/2$ if $|\epsilon| \leq \delta$, and $L(\epsilon) = \delta(|\epsilon| - \delta/2)$ otherwise. In Figure 4, $\delta = 2$.

model are as follows:

$$q_{t+1} - q_t = B_W \; E(q_t, \{x_t^i, \text{all } i\}) + \epsilon_t, \tag{22}$$

where $E(q_t, \{x_t^i, \text{all } i\}) = \frac{1}{I}\sum_i e^i(q_t, x_t^i)$.

Notice the *absence of weighting* in computing the total excess demands, and the *difference in timing of holdings* when evaluating excess demands. Also, under Walrasian dynamics, the matrix B_W should be *diagonal* with strictly positive diagonal coefficients.

Rather than running two separate regressions, we test whether the Walrasian model provides additional explanatory power beyond the ABL model. We do so by including a security's own excess demand $E(q_t, \{r_t^i, \text{all } i\})$ as a regressor in the corresponding equation of the ABL model (21). To avoid issues of multicolinearity, we *orthogonalize* the regressors of the Walrasian model with respect to the ABL regressors.[29] We then test whether the coefficients of the orthogonalized Walrasian excess demands are significant and positive. If so, the Walrasian model is deemed to provide explanatory power for price changes beyond the ABL model. If the coefficients are insignificant or negative, we conclude that the Walrasian model either does not provide explanatory power beyond the ABL model or makes the wrong predictions.

Allocations

The equations in (17) specify the evolution of holdings of risky assets under the ABL model. We set $\Delta = 1$, as for price dynamics, and assume equal impatience parameters ($c^i = \bar{c}$). The

[29]Orthogonalization is implemented by taking the difference between E and WE. Inspection of the resulting regressors reveals that the orthogonalized regressors equals the differences between the risk-aversion weighted per-capita holdings of a security and the unweighted per-capita holdings. The latter equals the number of units of the security in the market portfolio, i.e., the corresponding element in $\bar{r}$. Orthogonalization has at least one important effect. While Walrasian aggregate excess demands are not affected by the distribution of holdings across participants with different risk aversion, the orthogonalized Walrasian aggregate excess demands *are*, since the regressors in the orthogonalization, the ABL regressors, change with the distribution of holdings.

latter implies that the second term drops out. We are left with:

$$r_{t+1}^i - r_t^i = -\bar{c}\Omega \left(a^i r_t^i - \frac{\sum_i a^i r_t^i}{I} \right). \tag{23}$$

To interpret these equations, remember that i's willingness to pay is $\rho^i(x^i) = \mu - a^i \Omega r^i$. Therefore, (23) states that agents' allocations change in proportion to their willingness to pay relative to that of the average agent. This translates into the following predictions.

1. Allocations, scaled for risk aversion, change depending on *how far an agent's current holdings deviate from per-capita holdings, scaled for risk aversion.*

2. *Cross-security effects*: if holdings in one security deviate from risk-aversion scaled per-capita holdings, then this affects subsequent changes in holdings of *other* securities. The effects depend on payoff covariances.

As to the second point, if an agent holds too much of a security (scaled for risk aversion) relative to the risk-aversion weighted average holdings, and another security has payoffs with positive correlation, the agent will *reduce* holdings of the other security as well.

The scaling of an agent's holdings by risk aversion has its origin in the fact that a risk averse agent (an agent with high a^i) will always invest less in risky securities. Portfolio separation predicts how much less: the ratio of investments in a risky security of an agent relative to the average agent is described entirely by the ratio of the agent's risk aversion coefficient and the average risk aversion coefficient.[30] As such, portfolio separation is the crucial driver of allocation dynamics in the ABL model.

Risk-aversion scaled per-capita holdings provide a crucial benchmark in ABL allocation

[30] This can readily be derived from Equation (18).

dynamics. Because of its importance we shall refer to them with an acronym: RASE, for Risk-Aversion Scaled Endowment portfolio. The number of units RASE invests in each security are collected in the vector $\frac{\sum_i a^i r^i}{I}$. Compare this to the market portfolio, which in general features different investments: $\frac{\sum_i r^i}{I}$. We discuss later that the RASE portfolio provides predictions for the cross-section of prices of risky securities that are analogous to those of the market portfolio. The difference is that the predictions of the RASE portfolio hold off equilibrium as well. The market portfolio makes valid predictions only in equilibrium.

By adding error terms to (23), we translate the equations into regressions that we can bring to the data:

$$r^i_{t+1} - r^i_t = B \text{ WDeltaRASE}(t) + \epsilon_t, \tag{24}$$

where $\text{WDeltaRASE}(t) = \Omega\left(a^i r^i_t - \frac{\sum_i a^i r^i}{I}\right)$.

Tests focus on the elements of the coefficient matrix B. The matrix should be diagonal, with *strictly negative* diagonal elements. From an econometric point of view, however, the regression in (24) is problematic. Figure 3 displayed the evolution of deviations of average holdings of a group of participants from a benchmark (the market portfolio). The figure shows that the deviations are highly persistent. We expect this persistence to emerge in the regressors in (24) as well. Specifically, we expect the dynamics of the regressors to be close to unit-root. This induces huge autocorrelation in the error terms, which then causes significant biases in coefficient estimation, and mis-specification of standard errors. To avoid these issues, we therefore take first-differences. Investigation of the autocorrelation of error terms indicates that this was the right strategy.

We do not run allocation regressions on each participant separately. Instead, as we did for Figure 3, we average holdings across a homogeneous group of participants. A group is defined by (i) the risk aversion parameter of its members, and (ii) their initial allocations.

Setting Δ

We have set $\Delta = 1$. What does this mean practically? Is one (time) tick equal to one trade? Or, in calendar time, one second? The theory only assumes that Δ is long enough for everyone to trade, no matter how little. In practice, some participants trade only occasionally, and others trade a lot. Indivisibility makes it unprofitable for many to trade over very short intervals. As compromise, we measure time in terms of trades, not seconds, and take one time step to be equal to five trades. That is, $\Delta = 5$ trades. This is rather arbitrary, but reflects our intent to minimize bias while retaining power.[31]

3.3 Experimental Design

We report results from nine *sessions* with two risky securities and one risk-free security. Like cash in the experiments, the risk-free security did not earn interest. Because it could be sold short, it allowed participants to borrow money, at an interest set by the market.[32]

The first four sessions entailed two sets of four periods (for a total of eight). *Treatments* were distinguished by the sign of the covariances between the payoffs of the risky securities. Within a treatment, the four (4) *periods* were identical and independent replications, starting with the same initial endowments and the same mean-variance payoff functions. There were three groups of participants: one with a high coefficient of risk aversion, the other two with low coefficients of risk aversion. Table 1 lists the parameters of the four sessions. The table

[31]Shorter time intervals lead to biases towards finding no effect from the regressors, and longer time intervals cause lack of power because of reduced data points. We ran robustness tests and found the inference to be unchanged when Δ was set to 10 trades; power was reduced, however.

[32]For readers unfamiliar with markets experiments, Online Appendix F.1 briefly explains how they are run.

also reports CAPM equilibrium price predictions.

In the last five sessions, the sign and magnitude of payoff covariances were fixed for three periods. Hence, there were two treatments of three periods each. In contrast to the earlier sessions, initial endowments varied across the three periods within a treatment. As a result, CAPM equilibrium price predictions changed across *all* periods. Participants were divided into two groups depending on their coefficient of risk aversion (high; low).

Table 1 lists the parameters for the first four sessions.[33] Corresponding CAPM equilibrium price predictions are included as well.[34] Trade took place in online, anonymous, continuous open book systems. These systems are an expanded version of the traditional CDA whereby inferior limit orders are kept in an open book, until executed, or until canceled. In the first four sessions, the online system was Marketscape, developed by Charles Plott at Caltech.[35] In the subsequent five sessions, the online trading system was Flex-E-Markets, the same system used for the class experiment discussed earlier. Flex-E-Markets was originally developed by Peter Bossaerts and Elena Asparouhova, and now augmented by Jan Nielsen. Flex-E-Markets is available for use as a Software as a Service (Saas) through adhocmarkets.com.[36]

Participants were given a color-coded look-up table that, for every combination of holdings of the two securities (A and B) indicated their performance (utility) excluding payoffs from holdings of risk-free securities ("Notes") and cash. See Online Appendix H for a full set of instructions.

Participants were not informed of performance schedules or initial holdings of others. This

[33] The parameters for the sessions 5-9 can be found in Online Appendix G Table OA.3.

[34] In some of the periods in the sessions listed in Table OA.3, exchange took place with a one-shot call market. We exclude those periods since our theory does not apply to this exchange mechanism.

[35] Marketscape was also used in, e.g., Asparouhova, Bossaerts and Plott (2003); Bossaerts, Plott and Zame (2007).

[36] Flex-E-Markets provided the trading interface for the experiments reported in, e.g., Asparouhova and Bossaerts (2017); Asparouhova e.a. (2016).

way, those with knowledge of general equilibrium theory could not possibly derive equilibrium prices. This also means that participants could not form reasonable expectations about where prices would tend to, rendering credibility to the assumption of Local Optimization (Hypothesis 4). The number of participants fluctuated between 18 and 41, which is high relative to other market experiments. Earlier studies have suggested that, with multiple simultaneous markets, more than the usual number of participants (8-10) are needed in order for general equilibrium to emerge convincingly (Bossaerts and Plott, 2004).

In Sessions 1-4, accounting was done in an experimental currency converted to dollars at the end of a session at a pre-announced exchange rate. In the remaining sessions, all accounting was done in U.S. cents. Sessions lasted approximately three hours and the average payoff was $45 (with range between $5 and $150). The experiments were approved by the Caltech and University of Utah Institutional Review Boards (ethics committees). Instructions with snapshots of the MarketScape and Flex-E-Markets trading interfaces can be found in Online Appendix H.

3.4 Statistical Analysis

We perform regression analysis based on equations (21), (22) and (24). To study the slope coefficients, we report z-statistics based on robust regressions using Huber's method, with outlier parameter (δ) equal to 2.0 throughout, as explained before.[37]

Our data consist of price and allocation records for 2 securities in 9 experimental sessions and 2 treatments within each session, for a total of 36 samples/time series. Rather than reporting 36 z-statistics separately for each sample,[38] we report the distribution of the 36

[37]We implement Huber's robust regression using the method "robustfit" of the Matlab statistics package.

[38]Example: there are 36 z-statistics that test whether the diagonals in the coefficient matrix of (21) equal zero. Another example: there are 36 z-statistics that test whether the slopes on the orthogonalized Walrasian

Session 011128	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 0 (14)	8	2	0.028
Type 1 (14)	2	8	0.015
Type 2 (13)	2	8	0.23
Securities:			
Market (Units)	4.05	5.95	
Exp Payoff ($)	2.30	2.00	
Pay Variance	1.0	0.14	
Periods 1-4:			
Pay Covariance	-0.3		
CAPM Price	2.24	2.01	
Periods 5-8:			
Pay Covariance	0.3		
CAPM Price	2.14	1.94	

Session 020320	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 0 (10)	8	2	0.028
Type 1 (10)	2	8	0.015
Type 2 (10)	2	8	0.23
Securities:			
Market (Units)	3	4	
Exp Payoff ($)	2.30	2.00	
Pay Variance	1.0	0.14	
Periods 1-4:			
Pay Covariance	0.3		
CAPM Price	2.14	1.94	
Periods 5-8:			
Pay Covariance	-0.3		
CAPM Price	2.24	2.01	

Session 020424	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 0 (13)	8	2	0.028
Type 1 (13)	2	8	0.015
Type 2 (14)	2	8	0.23
Securities:			
Market (Units)	3.95	6.05	
Exp Payoff ($)	2.30	2.00	
Pay Variance	1.0	0.14	
Periods 1-4:			
Pay Covariance	0.3		
CAPM Price	2.13	1.94	
Periods 5-8:			
Pay Covariance	-0.3		
CAPM Price	2.24	2.01	

Session 020528	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 0 (13)	8	2	0.028
Type 1 (13)	2	8	0.015
Type 2 (14)	2	8	0.23
Securities:			
Market (Units)	3.95	6.05	
Exp Payoff ($)	2.30	2.00	
Pay Variance	1.0	0.14	
Periods 1-4:			
Pay Covariance	0.3		
CAPM Price	2.13	1.94	
Periods 5-8:			
Pay Covariance	-0.3		
CAPM Price	2.24	2.01	

Table 1: Parameters: Session 1-4. An experimental currency was used, converted to U.S. dollars at a pre-announced exchange rates. All parameters are expressed for 100 units of experimental currency. Type 0, Type 1 , and Type 2 subjects all had initial allocation of 0 Notes and 4.0 of Cash.

z estimates. Under the null hypothesis that the corresponding parameter equals zero, the distribution should be $N(0,1)$ (standard normal). The ability to use the entire empirical distribution of statistics across multiple samples is a luxury that experimental replications afford. For an earlier implementation of this approach, see Bossaerts, Plott and Zame (2007).

Under the alternative hypothesis (when the slope is non-zero), the z-statistics should continue to be Gaussian with unit variance, but with non-zero mean. The sizes of the effects under the alternative hypothesis could vary from one outcome to another, being governed by cohort-specific parameters such as the impatience and liquidity parameters. Hence, under the alternative hypothesis, we expect that, across sessions/treatments, the z-statistics behave as a Gaussian random variable with a *random mean*. That is, the z-statistic is a mixing Gaussian random variable with mixing on the mean. This implies that the distribution will still be Gaussian, but with variance larger than 1. See Figure 5, Left Panel.

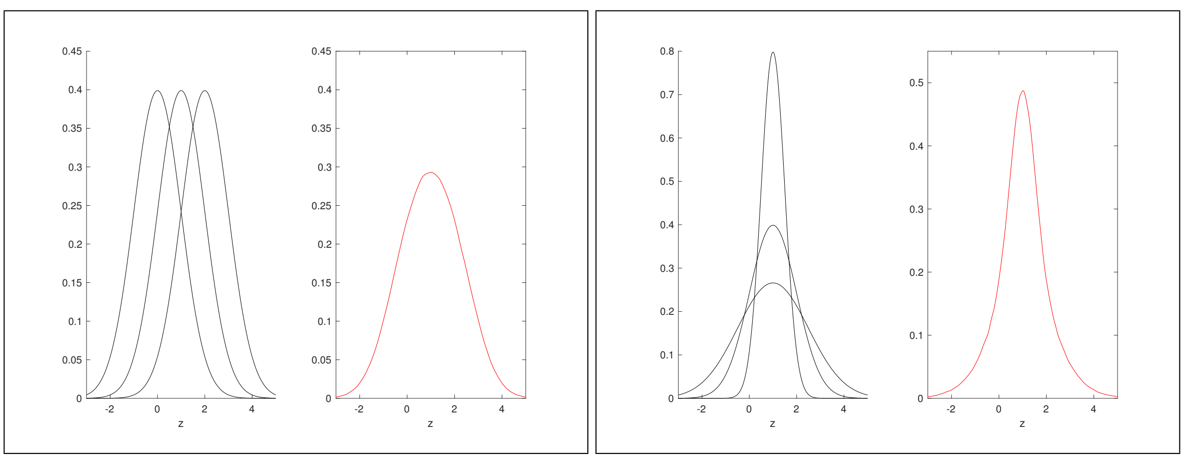

Figure 5: Left Panel: Under the ABL model, the true diagonal elements of the regression matrix B in (21) depend on parameters that are cohort and/or security specific, such as impatience and liquidity. As a result, while the corresponding z-statistics will still have unit variance (asymptotically), their mean changes randomly across the 36 time series. The unconditional distribution, shown to the right, will still be Gaussian, however. Right Panel: If z-statistics are computed using the wrong standard errors, and the standard errors are random across the 36 time series, the resulting unconditional distribution of z-statistics will be more peaked and exhibit heavier tails than the Gaussian distribution. Leptokurtosis therefore reveals model mis-specification.

excess demands [see (22)] provide no additional explanatory power for price changes beyond the regressors in (21).

The approach facilitates diagnostics on the correctness of the standard errors with which the z-statistics are constructed. If the standard errors are computed incorrectly, one could reasonably expect the z-statistics to be Gaussian with a standard deviation different from 1. The standard deviation may even depend on the sample (outcome) at hand. Consequently, the resulting distribution of z-statistics becomes a mixture-of-normals, with mixing on the standard deviation. This is well known to generate leptokurtosis: a density with excessive peaks and tails relative to the Gaussian distribution. See Figure 5, Right Panel. Consequently, leptokurtosis in the estimated density of the z-statistics will reveal mis-specification of the model with which standard errors are computed.

We estimate the density of the z-statistics using standard kernel smoothing techniques.[39]

4 Results

4.1 Prices

Diagonal Elements of B in (21). Figure 6a plots the 36 estimated z-statistics for the *diagonal* elements of the coefficient matrix in projections of price changes onto risk-aversion weighted excess demands pre-multiplied by the payoff covariance matrix. These are the diagonal elements of B in (21). There are 36 observations since there are 36 samples (time series), one for each of 2 assets per session-treatment, and for each of 18 session-treatments.[40] The 36 observations are depicted by stems on the horizontal axis of the plot. Under the null that the ABL model does not predict price changes, and provided the usual assumption for (asymptotic) gaussianity of the z-statistics is satisfied, the density of the z-statistics is $N(0, 1)$,

[39]We use the ksdensity method in the statistics package of Matlab.

[40]As mentioned before, a treatment consists of replications (periods) within a session with the same payoff covariance matrix but not necessarily the same initial allocations.

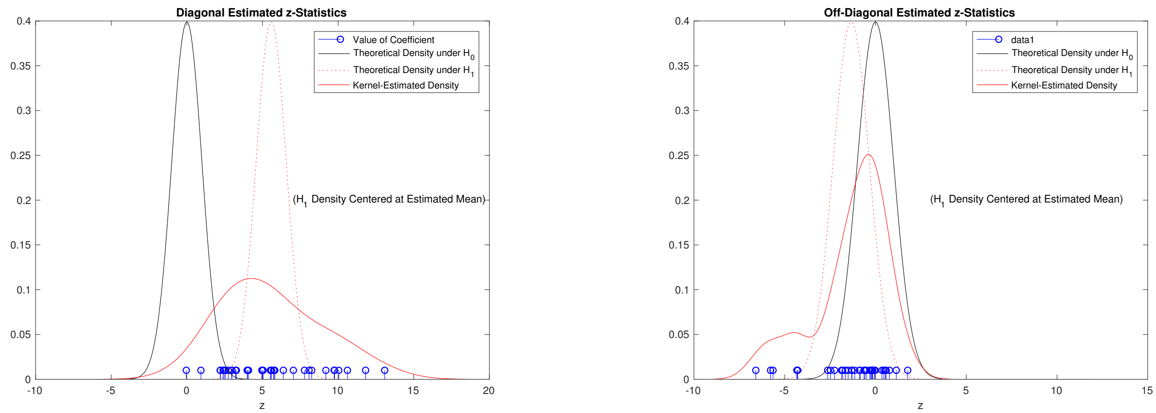

(a) Diagonal Coefficients

(b) Off-Diagonal Coefficients

Figure 6: Plot of 36 estimated z-statistics (stems) corresponding to diagonal (left) and off-diagonal (right) elements of the coefficient matrix B in (21). All transaction price changes over intervals of five (5) trades for a security in one session-treatment constitute a sample from which a single z-statistic is estimated. Solid red curve depicts kernel-estimated density of the z-statistics. Dotted red curve depicts Gaussian curve centered at the mean z-statistic and assuming unit variance; this is the theoretical curve under the alternative of a non-zero coefficient, centered at the observed mean, and assuming equal impatience and liquidity parameters across securities/sessions/treatments. Solid black curve depicts $N(0, 1)$, the theoretical density under the null that the coefficients are zero.

as indicated by the solid black curve. According to our theory, however, the diagonal elements of B should be strictly positive. As is clear from the figure, *all* 36 estimated z-statistics are positive. Their mean is indicated by the value of z where the red-dotted line reaches its peak. At more than 5, this mean is in the tails of the density of the z-statistics under the null, with a p value that is less than 10^{-6}. On these two accounts, we find strong confirmation of the theory.

The solid red line in Figure 6a displays the estimated density of the z-statistics. It is to be compared to the red dotted line, which represents the density centered at the mean z-statistic, and with variance equal to 1. This means that the red dotted line represents the distribution of the z-statistic under an alternative hypothesis whereby the true value of the diagonal coefficient is constant. The fact that the estimated density is flatter reveals that the true value of the diagonal coefficients varies across outcomes. This is not surprising since the true value depends on liquidity and impatience parameters which can be expected to vary across subject cohorts and securities. As a result, and if the standard errors were correctly specified, the true distribution of the z-statistic is Gaussian, with strictly positive mean. That is, the density should look like the red curve in Figure 5 of the Methods Section. Notice also that the estimated density (the solid red line) displays the typical bell shape of a Gaussian distribution. Disregarding slight positive skewness, the red curve in Figure 6a looks Gaussian.

Off-Diagonal Elements of B in (21). According to our theory, the *off-diagonal* elements of the coefficient matrix B in (21) should be zero. This reflects the fact that, once risk-aversion weighted excess demands are adjusted for the covariance matrix, cross-security effects should disappear. Figure 6b presents the evidence. The 36 estimated z-statistics of the off-diagonal coefficients are clearly clustered around zero, though there are a few large, negative outliers. The estimated density of the z-statistics (solid red line) overlaps substantially with the theoretical density under the null hypothesis (solid black line). The peak (mode) of the estimated density is close to zero (though negative). The mean estimated z-statistic, indicated by the peak of the dotted red density, is much further to the left, but still comfortably above -2 (the chance of observing an outcome of -2 or less under the null is approximately 2%). The outliers cause left-skewness in the density of the estimated z-statistics, which pushes the mean downward. With the exception of the negative skewness, the estimated density of the z-statistics (red curve) appears to be bell-shaped, suggesting that the z-statistics are well-specified.

Walrasian Dynamics: Diagonal Elements of B_W. We now turn to Walrasian influence on price dynamics. We determine to what extent price changes that are not captured by the

ABL model can be explained by traditional Walrasian excess demands. That is, we compute z-statistics for the diagonal elements of the coefficient matrix B_W in (22), after orthogonalizing the regressors with respect to the regressors in the ABL model (i.e., the regressors in (21)). Figure 7 displays the resulting 36 estimated z-statistics. They are mostly clustered around zero, consistent with the hypothesis that Walrasian dynamics cannot explain anything beyond Marshallian dynamics. There is one big (negative) outlier, beyond -5. The estimated density of the z-statistics (solid red curve) mostly coincides with the density under the null (solid black curve), though the left tail is a bit larger because of the outlier. The former has a mode close to 0, consistent with the null. If we look at the theoretical density centered at the sample mean z-statistic (dotted red curve), we observe that it is displaced to the left, which is again the influence of the outlier. We conclude that the preponderance of evidence points towards inability of Walrasian excess demands to provide explanatory power that is not already captured by the ABL model.

We emphasize that the negative outlier, and indeed all significantly negative outcomes, are inconsistent with Walrasian dynamics. If Walrasian dynamics truly explained some of the variance of price changes left unexplained by our theory, the test statistics should be *positive*. The vast majority are negative instead.

To better understand the meaning of the – often negative – z-statistics for the Walrasian excess demands, we plot them against (i) the estimated z-statistics corresponding to the diagonal elements of the coefficient matrix in the ABL model (the matrix B), and (ii) the estimated z-statistics corresponding to the off-diagonal elements of the same matrix.

Figure 8a plots the former. We observe a mild ($p = 0.05$) negative relationship. This means that, if we find a stronger positive influence of "driver" of a security's price according to the

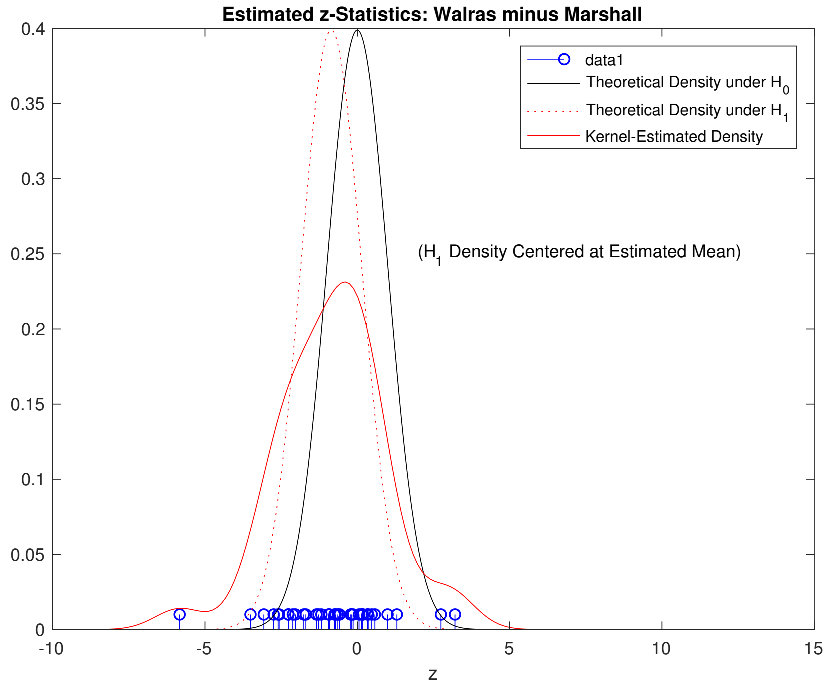

Figure 7: Plot of 36 estimated z-statistics (stems) corresponding to diagonal elements of the coefficient matrix B_W in (22); regressors are orthogonalized with respect to regressors in (21). See caption of Figure 6a for further information.

ABL model, we tend to find it offset by a negative influence of the security's own Walrasian excess demand.

But the latter has been *orthogonalized* with respect to the former. As mentioned before, the orthogonalized regressor equals to the difference between the risk-aversion weighted holdings of the security and the unweighted holdings (total supply). If risk averse subjects hold more of the security than others, the orthogonalized regressor is positive. Since its coefficient is negative, the induced price change is negative. It is intuitive what this is telling: risk averse agents pull down prices if they are holding too much of a risky security. Effectively, the ABL model under-estimates how much risk averse agents are willing to pull down prices. While we have been assuming that impatience is the same across agents, risk averse participants appear to be more impatient. This is consistent with subject-level data reported in Asparouhova and

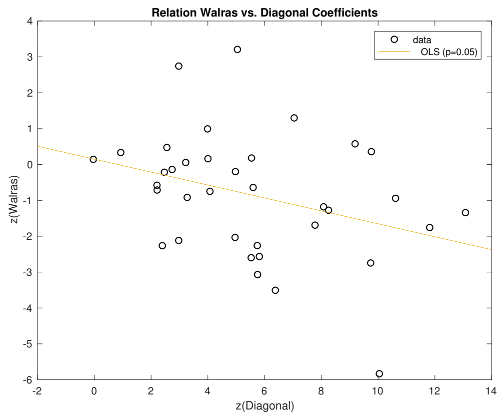

(a) Diagonal Elements of B

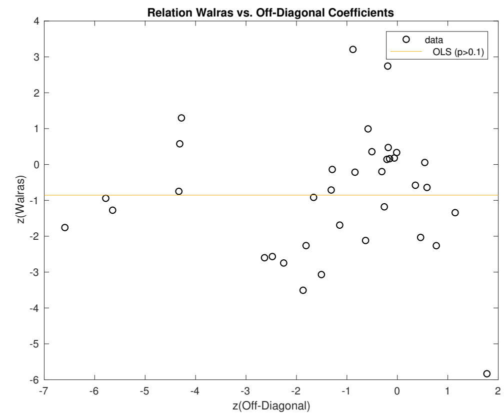

(b) Off-Diagonal Elements of B

Figure 8: Plot of relation of 36 estimated z-statistics corresponding to diagonal elements of the coefficient matrix B_W in (22) (regressors are orthogonalized with respect to regressors in (21)) and 36 estimated z-statistics corresponding to diagonal (left panel) and off-diagonal (right panel) elements of the coefficient matrix B in (21). All transaction price changes over intervals of five (5) trades for a security in one session-treatment constitute a sample from which a single z-statistic is estimated. The left panel's slope of the linear regression (yellow line) is significant at $p = 0.05$. It is insignificant ($p > 0.10$) on the right panel.

Bossaerts (2009).

No such relationship can be discerned when plotting estimated z-statistics for the orthogonalized Walrasian excess demands against the estimated z-statistics corresponding to the off-diagonal elements of B (point (ii) above). See Figure 8b.

By transforming the ABL regressors using Ω, we obtain an elegant way to compare data across treatments. Lost in this transformation is the *difference in dynamics* between the treatments: cross-security impact of excess demands on price changes are significant and *of opposite sign.* By merely changing the signs of the payoff covariances, we managed to induce fundamentally different price dynamics. See Asparouhova, Bossaerts and Plott (2003) for direct evidence, including experiments with three (rather than two) risky securities.

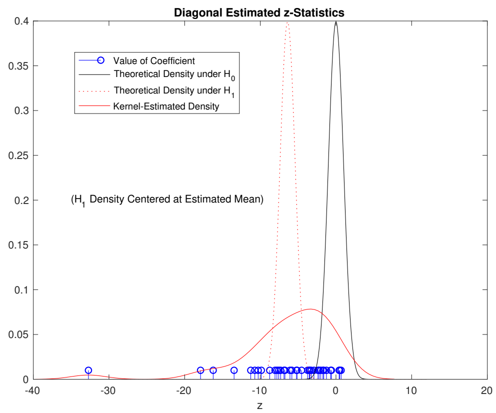

(a) Diagonal Elements of B

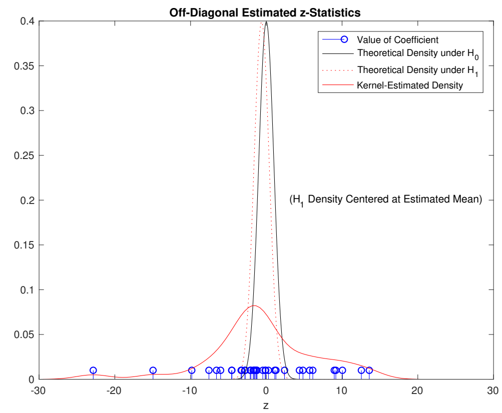

(b) Off-Diagonal Elements of B

Figure 9: Plot of 36 estimated z-statistics (stems) corresponding to diagonal (left panel) and off-diagononal (right panel) elements of the coefficient matrix B in (24). All allocation changes over intervals of five (5) trades for a security in one session-treatment constitute a sample from which a single z-statistic is estimated. Only allocation changes of the subject group with the highest risk aversion coefficient are included. Regressand is the average allocation change in that group. See the caption of Figure 6a for further information.

4.2 Allocations

Diagonal Elements of B in (24). Figure 9a displays the z-statistics pertaining to the diagonal elements of the coefficient matrix B in (24). These are z-statistics for the 36 security-session-treatment regressions of changes in per-capita holdings of the most risk averse subject group onto the difference in risk-aversion scaled holdings of the two securities and the per-capita risk-aversion scaled holdings, pre-multiplied by the payoff covariance matrix. The ABL model predicts *negative* coefficients. Figure 9a shows that, with a few exceptions, the z-statistics are indeed negative. The vast majority have values beyond the critical bound -2 (corresponding to $p = 0.02$). As before, the theoretical density of the z-statistics under the null that the regressor does not correlate with allocation changes is depicted with a solid black curve.

The solid red line depicts the estimated density of the z-statistics. Most of the mass is outside the interval of z-statistics where, under the null of no effect, 96% of the outcomes live, namely $[-2, 2]$. The dotted red line indicates the theoretical density of the z-statistics under the alternative that the effect is the same as that for the average z-statistic. This density hardly overlaps with that under the null hypothesis. The estimated density of the z-statistics is far more spread out, however, suggesting that the diagonal coefficients differ across session-treatments and securities. We presume that the heterogeneity emerges because of differing impatience and/or liquidity parameters. Ignoring the outlier, the estimated density is bell-shaped, suggesting that the standard errors are well-specified.

Overall, these statistical results provide strong support for our theory.

Off-Diagonal Elements of B in 24. Off-Diagonal elements of the regression coefficient matrix B should be zero in (24). Figure 9b shows that the z-statistics straddle zero, and that the theoretical density centered around the mean estimate (dotted red line) overlaps largely with the theoretical density under the null hypothesis of no effect (solid black line). However, the estimated density of the z-statistics (solid red line) is far more spread out than that under the null. This suggests that the true coefficients could be random, with a mean indistinguishable from zero. That is, our theory works *on average*, but there are deviations that the theory cannot explain. Evidently, these deviations can go either way.

We have defined a "treatment" as a sequence of periods in a session where the payoff covariance matrix is kept positive. In Sessions 5–9, initial allocations, and hence, equilibrium prices, changed across periods in a treatment. However, in Sessions 1–4, everything else remained the same across the periods of a treatment. As a result, the within-treatment periods were identical replications. Because of this, there is a possibility that participants started building expectations of price changes. Our theory assumes that agents cannot reasonably build expectations, and hence, behave in a myopic way (Hypothesis 4). We also tested the ABL model on a subset of unique periods within each treatment. Qualitatively, the inference is the same.[41]

5 Implications for Finance

5.1 Asset Pricing

Financial economists are interested in models that relate asset prices to covariances of their payoffs with some measure of aggregate risk. The CAPM provided the first example of this type of model. There, the price of an asset decreases in the covariance between its payoff and the payoff on the market portfolio. The market portfolio contains all risky securities, with units assigned to each security equal to the per-capita endowments. Roll has shown that CAPM obtains because, in equilibrium, the market portfolio is mean-variance optimal (Roll, 1977).

Here, we identify a portfolio with which to price all risky securities *even off-equilibrium*. We search for a benchmark portfolio that is mean-variance optimal throughout equilibration. Roll (1977) has shown that such a portfolio always exists (barring arbitrage opportunities), and that it prices all assets as follows. For a mean-variance efficient portfolio with composition (vector of units of each of the assets) v, there exists a scalar $\beta > 0$ so that:

$$q = \mu - \beta \Omega v. \tag{25}$$

Recall that, at any moment during the ABL equilibration process, prices follow a system of difference equations that depend on the weighted averages of agents' marginal rates of substitution. See (10). This system of difference equations pushes prices towards levels where

[41] Results are available upon request, and will be posted online together with the dataset and the statistical programs.

they are equal to those averages: $q \to \bar{\rho}$. Translated to our economy with quasi-linear preferences, and assuming that impatience parameters are equal across agents, this implies that prices exponentially converge to $q^* = \bar{\rho}(x) = \mu - \Omega \frac{1}{I} \sum a^i r^i$. The interpretation of this system of equations becomes clearer if we re-write it as follows:

$$q^* = \mu - \beta\Omega \sum_{i=1}^{I} \frac{a^i}{\sum_j a^j} r^i, \text{where } \beta = \frac{\sum_j a^j}{I}. \tag{26}$$

With reference to (25), this means that prices tend to make a particular portfolio mean-variance optimal. The portfolio is the one constructed from weighing holdings (r^i) with risk aversion (a^i). We referred to this portfolio before as the Risk-Aversion Weighted Endowment Portfolio (RASE). Mathematically, the weights equal $\frac{a^i}{\sum_j a^j}$.

We thus have obtained the remarkable result that, *throughout equilibration, prices tend towards levels that make the RASE portfolio mean-variance optimal.* Even if the market portfolio is off the mean-variance frontier throughout, RASE will tend towards it. Of course, as agents trade, their portfolios of risky assets will gradually converge (weights will become the same), while they will generally end up with different holdings of the numeraire. The RASE portfolio eventually converges to the market portfolio.

Because the result is only a tendency,[42] we cannot claim that the RASE portfolio is mean-variance optimal throughout equilibration. Instead, we make a weaker prediction, which is that the Sharpe ratio of the RASE portfolio is continuously higher than that of the market portfolio. The Sharpe ratio of a portfolio is the ratio of expected return in excess of the risk free rate and the return volatility. The return is defined as the end-of period payoff of the

[42] We qualified the result as a *tendency.* (26) is the steady-state point of a dynamic set of equations for prices. If allocations change before reaching the steady state, the dynamic set changes. The *nature* of the steady-state point does not change, however: it remains the risk-aversion weighted endowment portfolio. Still, the *weights* change, however, as holdings shift through trade.

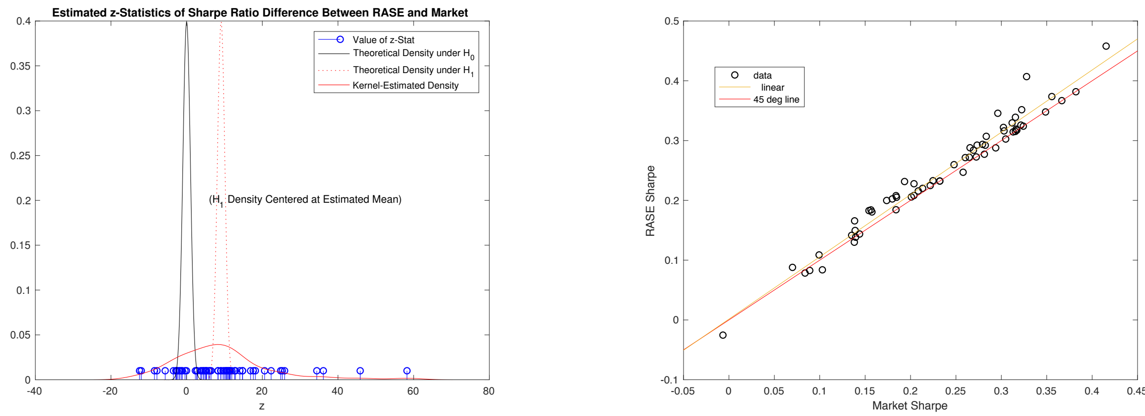

(a) Sharpe Ratio Test z-statistics

(b) Off-Diagonal Elements of B

Figure 10: Panel (a) plots of 62 estimated z-statistics (stems) testing whether the average Sharpe ratio of the RASE portfolio is higher than that of the market portfolio. Sharpe ratios are re-computed every 5 trades. Each period in the experiment generates one sample for which a z-statistic is estimated. Solid red curve depicts kernel-estimated density of the z-statistics. Dotted red curve depicts Gaussian curve centered at the mean z-statistic and assuming unit variance. Solid black curve depicts $N(0,1)$, the theoretical density under the null that the Sharpe ratio differences are zero. Panel (b) plots the 62 average Sharpe ratios of the RASE portfolio against those of the market portfolio. Red line denotes 45 degree line. Orange line depicts best linear fit (slope is significantly larger than 45 degrees at $p = 0.09$).

portfolio divided by the value of the portfolio at most recent transaction prices. The Sharpe ratio is maximal for a mean-variance optimal portfolio. We test whether RASE has a higher Sharpe ratio than the market portfolio.

There are 62 periods across all 9 sessions.[43] We test whether the Sharpe ratio of RASE is higher than that of the market in these 62 periods. At intervals of 5 trades, we compute the RASE portfolio and evaluate its Sharpe ratio. We do the same for the market portfolio, and compute the difference between the Sharpe ratio of RASE and the market portfolio. We then calculate the average of this difference for the period, and the corresponding z-statistic. We thus obtain 62 z-statistics. Figure 10a plots them as (blue) stems.

The vast majority of the z-statistics are positive, and 46 out of 62 reach a value above

[43]Four (4) sessions with 8 periods and five (5) sessions with 6 periods.

2 ($p << 0.001$), confirming that the RASE portfolio tends to dominate the market portfolio in mean-variance space. However, there are quite a few negative observations as well, some of them way in the left tail of the theoretical density under the null that the two portfolios generate the same Sharpe ratio (black solid line). A comparison of the theoretical density under the alternative that the expected z-statistic equals the sample average[44] (dotted red line), and the estimated density of the z-statistics (solid red line) reveals substantial heterogeneity across periods. The former hardly overlaps with the theoretical density under the null, but the latter has a significant overlap in the left tail.

Sharpe ratios were re-evaluated every five (5) trades. It may be that intervals of five trades are insufficiently long for the hypothesized effect to emerge. But the tendency is apparent: there is a portfolio that most likely will generate a higher Sharpe ratio than the market portfolio.

In historical data from field stock markets, the Sharpe ratios of proxies of the market portfolio have been found to be lower than those of portfolios that put more weight on, say, high-value stock and smaller firms. See Fama and French (2004). It would be interesting to determine to which extent the weight adjustments needed to beat market proxies in the field reflect differences in holdings of component securities across investors with varying levels of risk aversion. These adjustments make up the differences between RASE and the market portfolio. Of course, the lower performance of the market proxies in historical data from the field may also reflect that these are only proxies, and not the true market. In our experiments, we know what the true market portfolio is. Regardless, the finding that RASE tends to dominate in terms of Sharpe ratio even off-equilibrium provides a sensible alternative explanation for the poor historical performance of the market portfolio. We leave these and related issues for future work.

[44] Sample average of z-statistics $= 9.114$.

Further analysis of our data reveals that the RASE portfolio tends to perform better (in terms of Sharpe ratio) when the market Sharpe ratio is higher. Figure 10b plots the 62 average Sharpe ratios of RASE against those of the market portfolio. The solid red line depicts the 45 degree line. If an observation lies above this line, it implies that RASE performs better than the market portfolio. The dotted red line depicts a linear (OLS) fit (slope: 1.0416, $p = 0.09$ for null hypothesis that slope equals 45 degrees). The difference between the linear fit and the 45 degree line increases as the Sharpe ratio of the market increases: RASE tends to outperform more when the market portfolio generates a higher Sharpe ratio.

5.2 Momentum, Volume and Liquidity

Momentum.

In the ABL model, prices change in reaction to average marginal rates of substitution, see (10). Agents' marginal rates of substitution change in response to changes in holdings due to trade. As a result, a rich pattern of price dynamics is possible. In particular, it generates interesting cross-autocorrelations that, like the cross-security effects of risk-aversion weighted excess demands on price changes, depend on payoff covariances. Cross-autocorrelation intensities depend crucially on adjustment parameters, such as the liquidity parameters α_k and the impatience parameters c^i. This means that cross-autocorrelation patterns could provide statistical input to infer those adjustment parameters.

Interestingly, cross-autocorrelations have been recorded in historical field data. Importantly, they are thought to be the key factor behind the momentum effect, i.e., the finding that recent winners outperform recent losers, even after adjusting for risk (Lewellen, 2002). Momentum has always been considered to be puzzling. Here, momentum emerges as a feature of off-equilibrium dynamics, through cross-autocorrelations tied to adjustment dynamics. Indeed, prices of some securities adjust faster than others, because trade in those securities leads to larger utility increases, or because agents with higher risk aversion or trading impatience are disproportionately invested in them.

When analyzing the experimental data, however, we uncovered little evidence of momentum. Presumably, this is because, with only 2 risky securities, the power to discover momentum is reduced. We leave exploration of momentum in experiments with larger cross-sections for future work.

Remark 6. *Absent knowledge of economy-wide parameters, agents cannot exploit the features of price dynamics reported in the Results section. For instance, agents lack the information needed to form estimates of risk-aversion weighted excess demands, which are needed to predict price changes. Momentum, however, is a portfolio that can be constructed in the absence of structural knowledge of the economy. Since momentum should be profitable in our setting, some agents may want to exploit it. An interesting issue for future research is to determine to what extent this would cause equilibrium convergence to fail.*

Volume and Liquidity.

Our allocation dynamics have immediate consequences for volume, and hence, liquidity. To see how, remember individual allocation dynamics (23):

$$r_{t+1}^i - r_t^i = -\bar{c}\Omega \left(a^i r_t^i - \frac{\sum_i a^i r_t^i}{I} \right).$$

Now consider the following cases.

- *Case 1.* Everyone starts from the same initial allocations, meaning that all agents hold the market portfolio: $r_0^i = \bar{r}$. Risk aversion coefficients (a^i) are different, however. In this case, the *initial* adjustment is as follows:

$$r_1^i - r_0^i = -\bar{c}\left(a^i - \frac{\sum_i a^i}{I}\right)\Omega\bar{r}.$$

The changes in holdings are a linear transformation of the market portfolio. Except in the unlikely event that the market portfolio is an eigenvector of Ω, agents must initially trade away from the market portfolio. That is, they start from CAPM equilibrium holdings, *only to immediately deviate*. The more extreme one's risk aversion (a^i) is relative to the average, the farther away the initial movement is. Ignoring off-diagonal terms of Ω, the more risk averse agents sell securities, *focusing on the most risky ones* (highest variance). Likewise, less risk averse agents do what is locally optimal: increase risk exposure by prioritizing purchases of the most risky securities.

The effect of the off-diagonal elements of Ω, the payoff covariances, merits separate discussion. When the covariances are negative, agents' portfolios remain closer to the market portfolio than in the scenario when payoff covariances are zero or positive. The intuition is simple: when payoff covariances are negative, assets are natural hedges for one another. Increasing one's risk exposure by buying the most risky securities leads to a less diversified portfolio, i.e., to utility losses. Maximum local gains in utility are obtained by trading combinations of securities that are closer to the per-capita average endowment, i.e., the market portfolio. As a consequence, throughout equilibration, agents stay closer to the market portfolio than in the scenario where payoff covariances are zero or positive.

- *Case 2.* Agents start with different endowments but have the same risk aversion $a^i = \bar{a}$.

Then:

$$r_1^i - r_0^i = -\bar{c}\bar{a}\,\Omega\left(r_0^i - \bar{r}\right).$$

Here, agents adjust smoothly towards the market portfolio. Since the covariance matrix Ω multiplies the deviations of initial holdings from the market portfolio, adjustment will again be faster in the high-variance securities. As in Case 1, this effect will be attenuated if the off-diagonal elements of Ω (covariances) are negative.

The two cases reveal that adjustment will be faster in the high-variance assets. This means that *liquidity will initially be highest in the high-variance assets*. Negative off-diagonal terms (negative covariances) may partially offset this tendency.

But this only concerns liquidity when allocations are far from equilibrium. Closer to equilibrium, all efforts are concentrated on trading towards the market portfolio. In Case 1 above, individual holdings moved away from the market portfolio. Because the low-variance asset holdings have not been adjusted commensurate with the high-variance asset holdings, final adjustments are needed in the former, and hence, liquidity moves towards the low-variance assets when the economy is closer to reaching equilibrium allocations.

This is a novel prediction of our theory, worthy of further exploration, both in follow-up experiments with more than 2 risky assets, and in historical data from field markets.

A recent explanation of volume and liquidity has focused on optimal attention, see Galai and Sade (2006), followed by Karlsson, Loewenstein, and Seppi (2009) and Andries and Haddad (2020). There is a relationship between the explanations provided by these papers and ours. Agents' trade intensities are determined by the gradient of their utilities: agents trade faster in assets that provide a higher increase in utility. In optimal attention models, trade is also determined by assets that generate the highest potential change in utility.

Another recent theory of volume and liquidity has focused on portfolio separation; see Lo and Wang (2000). The reasoning is as follows. Since optimal portfolios can be described in terms of a limited number of benchmark portfolios, agents merely need to trade those portfolios. Absent direct access to the benchmark portfolios, trade in individual assets should only take place in proportion to the weights of the assets in the benchmark portfolios. Consequently, turnover (volume divided by total supply) is predicted to be constant across assets. As an example, take Case 1 above: all agents have the same endowment (hence, all endowments are a fixed combination of the riskfree asset and the market portfolio), but exhibit differing risk preferences. In the world of Lo and Wang (2000), more risk averse agents reduce their exposure to risk by trading the market portfolio with less risk averse agents, or, absent direct trade in the market portfolio, they trade the component assets in proportion to their weights in the market portfolio. As a result, volume will be proportional to weights in the market portfolio, and turnover will be constant.

Our predictions are markedly *different*: agents initially trade the asset with the highest variance (ignoring payoff covariances, which may attenuate the variance effect). Volume will therefore be proportional to (payoff) variance. As the economy approaches its equilibrium allocations, however, more trade will take place in the low-variance assets. Consequently, the relation between volume and variance is obscured by how far the economy is off equilibrium. One could turn this around: the relation between volume and variance is an indication of how far the economy is from equilibrium.

Interestingly, Lo and Wang (2000) show that, historically, volume on the NYSE and AMEX tends to increase when idiosyncratic risk is higher. Since idiosyncratic risk is a large proportion of total risk, this suggests that volume increases in (total) variance, consistent with *an economy* *that is far from equilibrium.*

6 Concluding Comments

Previous research has shown that standard global tatonnement and non-tatonnement are not consistent with within-period price dynamics in continuous double auctions (CDAs). Building on earlier experimental evidence from single CDAs (Friedman, 1991; Plott, 2000), we describe a Marshallian theory of the forces driving the economy to equilibrium. The theory is applicable to multiple, simultaneous CDAs and consistent with experimental findings with continuous double auction markets. Our theory was built from the level of the agents up, to obtain implications for market-wide price and allocation dynamics. Our theory is based on three main assumptions. One, agents in CDAs only submit (small) orders that maximize local gains from trade. Two, quantity moves to agents who offer the higher surplus to the market. Three, agents' bids are benchmarked against lagged prices.

In our experiments, we induced quasi-linear, mean-variance preferences in a way that makes the economy isomorphic to a CAPM one. The findings are in line with the theoretical predictions. Price changes correlated not only with own risk-aversion weighted excess demand, but also with risk-aversion weighted excess demands in other assets, in ways that related to the payoff covariance matrix. Traditional Walrasian excess demands either did not provide additional explanatory power or predicted price changes in a direction that is opposite to that expected. Our model correctly captured dynamics of the average allocation of participants stratified by risk aversion. Cross-equation effects emerged here as well, again determined by the covariance matrix (Hessian of the utility function).

Beyond price and allocation dynamics, we discovered that prices tend in a direction that makes one portfolio mean-variance optimal throughout equilibration. This portfolio, the risk-aversion scaled endowment portfolio (RASE), re-assigns weights in the market portfolio depending on the risk aversion of the agents holding the component assets.

Our results are not isolated to the experiments reported here. In Appendix F.2, we corroborate the findings in about 3200 observations from three sessions of four-asset experiments. The sessions differ from the ones reported on in the paper, in that: (i) mean-variance preferences are not induced; instead states are actually realized, though mean-variance preferences appear to capture price behavior well; (ii) there is no deliberate attempt to control the relation between excess demands and transaction price changes through changes in payoff covariances. In addition, Asparouhova, Bossaerts and Plott (2003) reports an analogous link between payoff covariances, on the one hand, and the relation between excess demands and price changes, on the other hand, in over 11,000 transaction price changes from eight sessions with three assets. Finally, Gillen e.a. (2021) also reports cross-security effects in price changes and excess demands in an unrelated, three-commodity experiment. Our theory generically predicts such cross-effects.

Much remains to be done. We have not allowed for speculation, and information (about final payoffs) was homogeneous. As to historical analysis of field markets, however, our findings should invite empiricists to re-assess prices, momentum, volume and liquidity, using our theory as guidance. One interesting question, for instance, is whether there is a relationship between our RASE portfolio and the factor portfolios that have historically out-performed buying and holding the market portfolio (in terms of Sharpe ratios).

References

ANDERSON, C., S. GRANAT, C. PLOTT AND K. SHIMOMURA (2004). Global Instability in Experimental General Equilibrium: The Scarf Example. *Journal of Economic Theory* 115:209–249.

ANDRIES, M. AND V. HADDAD (2020). Information Aversion. *Journal of Political Economy* 128:1901–1939.

ARROW, K. AND F. HAHN (1971). *General Competitive Analysis*. Holden-Day, San Francisco, CA.

ASPAROUHOVA, E. AND P. BOSSAERTS (2009). Modeling Price Pressure In Financial Markets. *Journal of Economic Behavior and Organization* 72:119–130.

ASPAROUHOVA, E. AND P. BOSSAERTS (2017). Experiments on percolation of information in dark markets. *Economic Journal* 127:F518–F544.

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.

ASPAROUHOVA, E., P. BOSSAERTS, N. ROY AND W. ZAME (2016). "Lucas" in the Laboratory. *Journal of Finance* 71:2727–2780.

AVELLANEDA, M., J. REED AND S. STOIKOV (2011). Forecasting prices from Level-I quotes in the presence of hidden liquidity. *Algorithmic Finance* 1:35–43.

BIAIS, B., MARIOTTI, T., MOINAS, S., AND S. POUGET (2017). Asset pricing and risk sharing in a complete market: An experimental investigation. Toulouse School of Economics (TSE) Working Paper 17-798.

Bossaerts, P. and C. Plott (2004). Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets. *Review of Finance* 8:135–169.

Bossaerts, P. and C. Plott (2008). From Market Jaws to The Newton Method: The Geometry of How a Market Can Solve Systems of Equations, in: *Handbook of Experimental Economics Results, Volume I.* Charles R. Plott and Vernon L. Smith (eds.). North-Holland, Amsterdam, 22–24.

Bossaerts, P., C. Plott and W. Zame (2007). Prices and Portfolio Choices in Financial Markets: Theory, Econometrics and Experiment. *Econometrica* 75:993–1038.

Budish, E., P. Cramton, and J. Shim (2015). The High-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130:1547–1621.

Champsaur, P. and B. Cornet(1990). Walrasian Exchange Processes, in: *Economic Decision Making: Games, Econometrics, and Optimisation*. J. Gabszewicz, J.-F. Richard and L. Wolsey (eds.). Elsevier, Amsterdam, 1–13.

Constantinides, G. M. (1990). Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy* 98:519–543.

Crockett, S. (2013) Price Dynamics in General Equilibrium Experiments. *Journal of Economic Surveys* 27:421–438.

Crockett, S., D. Friedman, and R. Oprea. (2017) Aggregation and convergence in experimental general equilibrium economies constructed from naturally occurring preferences. No. SP II 2017-501. WZB Discussion Paper, 2017.

Fama, E.F. and K.R. French (2004). The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*, 18:25–46.

Friedman, D. (1979). Money-Mediated Disequilibrium Processes in a Pure Exchange Economy. *Journal of Mathematical Economics* 6:149–167.

Friedman, D. (1991). A Simple Testable Model of Double Auction Markets. *Journal of Economic Behavior and Organization* 1:47–70.

Galai, D. and O. Sade (2006). The "Ostrich Effect" and the Relationship between the Liquidity and the Yields of Financial Assets. *The Journal of Business* 79:2741–2759.

Gillen, B., Hirota, M., Hsu, M., Plott, C. and Rogers, B. (2021). Divergence and Convergence in Scarf Cycle Environments: Experiments and Predictability in the Dynamics of General Equilibrium Systems *Economic Theory* 71:1033–1084.

Hahn, F. H., and T. Negishi (1962). A Theorem on Non-Tatonnement Stability. *Econometrica* 30:463–469.

Inoua, S. M. and V. L. Smith (2020). Classical versus Neoclassical Equilibrium Discovery Processes in Market Supply and Demand Theory. *Chapman University Economic Science Institute*, working paper.

Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty* 38:95–115.

Ledyard, J. O. (1974). Decentralized Disequilibrium Trading and Price Formation. Working paper #86 of The Center for Mathematical Studies in Economics and Management Science, Northwestern University.

Lewellen, J. (2002). Momentum and Autocorrelation in Stock Returns. *Review of Financial Studies* 15:533–563.

Lo, A. W., and J. Wang (2000). Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory. *Review of Financial Studies* 13:257–300.

Marshall, A. (1890). *Principles of Economics.* vol. I, London: Macmillan.

Negishi, T. (1962). The Stability of a Competitive Economy: A Survey Article. *Econometrica* 30:635–669.

O'Hara, M., C. Yao and M. Ye (2014). What's Not There: Odd Lots and Market Data. *The Journal of Finance* 69:2199–2236.

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

Plott, C. R. (2001) Equilibrium, Equilibration, Information and Multiple Markets: From Basic Science to Institutional Design. In Nobel Symposium on Behavioral and Experimental Economics, Stockholm, Sweden.

Plott, C.R. and K. Pogorelskiy (2017). Call market experiments: Efficiency and price discovery through multiple calls and emergent Newton adjustments. *American Economic Journal: Microeconomics*, 9:1–41.

Roberts, D.J. (1979). Strategic Behavior in the MDP Procedure. *Aggregation and Revelation of Preferences* J.J.Laffont (ed.), North Holland, 353–362.

Roll, R. (1977). A Critique of the Asset Pricing Theory's Tests. *Journal of Financial Economics* 4:129–176.

Rostek M. and M. Weretka (2015). Dynamic Thin Markets. *The Review of Financial Studies* 28:2946-2992.

Samuelson, P. (1947). *Foundations of Economic Analysis.* Harvard University Press, Cambridge, USA.

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

Smith, V. L. (1965). Experimental auction markets and the Walrasian hypothesis. *Journal of Political Economy* **73** 387–393.

Sundaresan, S. (1989). lntertemporally Dependent Preferences and the Volatility of Consumption and Wealth *The Review of Financial Studies,* 2:73–89.

Uzawa, H. (1962). On the Stability of Edgeworth's Barter Process. *International Economic Review* 3:218–232.

ONLINE APPENDIX

A Proofs

A.1 Incentive Compatibility of Optimal Bidding Strategy

About Remark 1: We prove here that, if both the quantity adjustment and the price setting rules are known, if $\alpha_k = \alpha, \forall k$, and *if bids are a local Nash equilibrium, then Hypothesis 4 is satisfied.*

Proof: Suppose all i believe Hypothesis 1; that is, $\Delta r_t^i = A(b_t^i - q_t)$. Further suppose they believe, as implied by Hypotheses 1 and 3, that $q_t = (1/I)\sum b_t^i$. Further suppose they choose b_t^i to be a local Nash Equilibrium. That is, for every i,

$$b_t^i \in \ \text{argmax}\ \Delta u_t^i = (\rho_t^i - q_t)A(b_t^i - q_t) \tag{OA.1}$$

$$= (\rho_t^i - \frac{\sum_j b_t^j}{I})A(b_t^i - \frac{\sum_j b_t^j}{I}) \tag{OA.2}$$

Letting $\bar{b}_t = \frac{\sum b_t^j}{I}$, the first order conditions for this are: $\frac{-1}{I}(b_{k,t}^i - \bar{b}_{k,t})\alpha_k + \frac{I-1}{I}(\rho_{k,t}^i - \bar{b}_{k,t})\alpha_k = 0$ or $b_t^i = \bar{b}_t + (I-1)(\rho_t^i - \bar{b}_t)$. Summing over i gives $\bar{b}_t = \hat{\rho}_t = \frac{\sum \rho_t^i}{I}$. So the local Nash equilibrium has $b_t^i = \hat{\rho}_t + (I-1)(\rho_t^i - \hat{\rho}_t)$. Since $q_t = \bar{b}_t = \hat{\rho}_t$ this means $b_t^i = q_t + (I-1)(\rho_t^i - q_t)$. Let $c^i = \alpha\frac{I-1}{\Delta}$.

A.2 Derivation of ABL Dynamics

From Hypotheses 1-4 , we have

$$r^i_{t+\Delta} - r^i_t = \Delta A(b^i_t - q_t) \tag{OA.3}$$

$$s^i_{t+\Delta} - s^i_t = -q_t \cdot (r^i_{t+\Delta} - r^i_t) \tag{OA.4}$$

$$\sum_i (r^i_{t+\Delta} - r^i_t) = 0 \tag{OA.5}$$

$$b^i_t - q_{t-\Delta} = c^i A^{-1}(\rho^i_t - q_{t-\Delta}) \tag{OA.6}$$

Substitute b_t from (OA.6) into (OA.3) to get

$$r^i_{t+\Delta} - r^i_t = \Delta A(q_{t-\Delta} - q_t) + \Delta c^i(\rho^i_t - q_{t-\Delta}) \tag{OA.7}$$

Sum (OA.7) over all i to get

$$q_t - q_{t-\Delta} = \Delta A^{-1}\bar{c}(\bar{\rho}_t - q_{t-\Delta}) \tag{OA.8}$$

Substitute (OA.8) into (OA.7) to get

$$r^i_{t+\Delta} - r^i_t = \Delta\left(-\bar{c}(\bar{\rho}_t - q_{t-\Delta}) + c^i(\rho^i_t - q_{t-\Delta})\right) \tag{OA.9}$$

A.3 Proof of (11)-(13)

The dynamics of our model are

$$
\begin{aligned}
q_{k,t} &= q_{k,t-\Delta} + \Delta \frac{\bar{c}}{\alpha_k}(\bar{\rho}_{k,t} - q_{k,t-\Delta}) \\
q_0 &= \bar{\rho}_0 \\
r^i_{k,t+\Delta} &= r^i_{k,t} + \alpha_k \Delta (q_{k,t-\Delta} - q_{k,t}) + c_i \Delta (\rho^i_{k,t} - q_{k,t-\Delta}) \\
s^i_{t+\Delta} &= s^i_t - q_t \cdot (r^i_t - r^i_{t-\Delta}).
\end{aligned}
$$

(The third equation uses (OA.7)).These contain a subtlety that must be dealt with if we want to let $\Delta \to 0$ to get the continuous version. This is a set of second-order difference equations since they specify dynamics over two intervals: $[t-\Delta, t)$ and $[t, t+\Delta)$. To get them into a standard set of first-order difference equations, let $z_t = q_{t-\Delta}$ and then, with a little algebra, rewrite the equations as:

$$
\begin{aligned}
z_{t+\Delta} - z_t &= \Delta A^{-1} \bar{c}(\bar{\rho} - z_t) \\
r^i_{t+\Delta} - r^i_t &= \Delta \left(c^i (\rho^i_t - z_t) - \bar{c}(\bar{\rho}_t - z_t) \right) \\
s^i_{t+\Delta} - s^i_t &= \Delta \left(-z_{t+\Delta} \right) \cdot \left(c^i (\rho^i_t - z_t) - \bar{c}(\bar{\rho}_t - z_t) \right)
\end{aligned}
$$

As $\Delta \to 0$, everything is well-behaved, and we end up with

$$
\begin{aligned}
\frac{dz}{dt} &= A^{-1} \bar{c}(\bar{\rho}_t - z_t) \\
\frac{dr^i}{dt} &= \left(c^i (\rho^i_t - z_t) - \bar{c}(\bar{\rho}_t - z_t) \right) \\
\frac{ds^i}{dt} &= -z_t \cdot \left(c^i (\rho^i_t - z_t) - \bar{c}(\bar{\rho}_t - z_t) \right)
\end{aligned}
$$

Now, note that as $\Delta \to 0, z_t = q_{t-\Delta t} \to q_t$. Substituting this, gives (11) - (13).

A.4 Proof of Theorem 1

Theorem 1: (Convergence to Pareto Optimality)

Let $x_t = (s_t, r_t)$*. If (i) there are no income effects, i.e.,* $u_0^i(x^i) = 1$ *(wlog) for all* i *and all* $x^i \in X$*, and (ii)* $x_t^i > 0$ *for all* t*, then for the dynamics in (8) and (10),* $(x_t, p_t) \to (x^*, p^*)$ *where* x^* *is Pareto-optimal and* $e(p^*, x^*) = 0$*.*

Proof: We use $\sum c^i u^i$ *as a Lyapunov function. Let* $\kappa^i = c^i(\rho^i - q)$*. Then we can write* $d(\sum_i c^i u^i)/dt = \sum_i c^i \frac{du^i}{dt} = \sum c^i(\rho^i - q)\frac{dr_t^i}{dt} = \sum_i c^i(\rho_q^i)[c^i(\rho^i - q) - \bar{c}(\bar{\rho} - q)] = [(\sum \kappa^i \kappa^i) - (1/I)\sum_k(\sum \kappa_k^i)(\sum \kappa_k^i)$*. By the triangle inequality,* $\sum ||\kappa^i||^2 \geq ||\sum \kappa^i||^2$*. Therefore* $\sum ||\kappa^i||^2 > (1/I)||\sum \kappa_k^i||^2$ *if* $\kappa^i \neq 0$ *for some* i*. Therefore,* $d(\sum c^i u^i)/dt > 0$ *unless* $\kappa^i = 0$ *for all* i *which is true iff* $\rho^i = q$ *for all* i*.* That is, convergence ends at a Pareto-optimal allocation.

Remark 7. *Condition (i) is included because we do not have a proof of convergence for utilities with income effects. We also do not have a counter example where such convergence will not occur. One could, of course, revise the model and impose a no-regret condition on trades that would ensure* $du_t^i/dt \geq 0$*. This would guarantee convergence. We do not do that here because, as we will see below, the model as it now stands is consistent with the data. If it is the right model of behavior in the CDA experiments, then a lack of convergence would be a feature and not a bug.*

Remark 8. *Condition (ii) is included above for technical reasons. If* $du_t^i/dt \geq 0$ *along the path for all* i*, then (ii) would not be necessary. But when* $du_t^i/dt < 0$ *is possible for some* i*, we need to worry about* x^i *hitting the boundary of the feasible consumption set. There are standard* *ways to modify (11)-(13) to deal with this. We do not pursue them here.*

B Local Marshallian Equilibrium (LME)

B.1 Theory

In the ABL model of individual behavior, Hypothesis 4, we assumed that bids at t are based on the prices and allocations arrived at in the interval $t - \Delta$. But another hypothesis might be that bids and prices are simultaneously determined within the time Δ. It is interesting to consider what the dynamics of price formation would then look like. We begin with

Hypothesis 6. *Local Optimization*

$$b_t^i = q_t + c^i \Delta A^{-1}(\rho^i(x_t^i) - q_t), \forall i, \forall t > 0.$$

It is easy to compute the local equilibrium in the interval $[t, t + \Delta)$.

Lemma 1. *Local Marshallian Equilibrium (LME). Under Hypotheses 1-3, and 6,*

$$q_t = \frac{\sum_i b_t^i}{I} = \bar{b}_t = \frac{\sum_i c^i \rho^i(x_t^i)}{\sum_i c^i} = \bar{\rho}(x_t).$$

Proof. Hypotheses 1-3 imply $q_t = \frac{\sum_i b_t^i}{I}$. Then summing b_t^i from Hypothesis 6 gives the desired result.

□

q_t is the local Marshallian equilibrium price, at which individuals will not want to change their bids and at which Marshallian trading is feasible.

The dynamics of the LME model are:[45]

$$r^i_{t+\Delta} = r^i_t + c^i\Delta(\rho^i(x^i_t) - q_t), \forall i, \tag{OA.10}$$

$$s^i_{t+\Delta} = s^i_t - q_t \cdot (r^i_{t+\Delta} - r^i_t),\ \forall i, \tag{OA.11}$$

$$q_t = \bar{\rho}(x_t). \tag{OA.12}$$

Dividing (OA.10) and (OA.11) by Δ and letting $\Delta \to 0$ leads to a continuous-time theory:[46]

$$\frac{dr^i_t}{dt} = c^i(\rho^i(x^i_t) - q_t),\ \forall i, \tag{OA.13}$$

$$\frac{ds^i_t}{dt} = -q_t \cdot \frac{dr^i_t}{dt},\ \forall i, \tag{OA.14}$$

$$q_t = \bar{\rho}(x_t). \tag{OA.15}$$

In continuous time, the process (OA.13)-(OA.15) will converge to a rest point from any initial price and allocation, *even if there are income effects*. This may not be true for (OA.10)-(OA.12) in discrete time if step sizes are too large.

Theorem 2. *(Convergence to Pareto Optimality)*

For the dynamics in (OA.13)-(OA.15), $(x_t, p_t) \to (x^*, p^*)$ where x^* is Pareto-optimal and (p^*, x^*) is a competitive equilibrium at x^*.

Proof. For each i, $\frac{du^i_t}{dt} = u^i_{0,t}(\rho^i_t - q_t) \cdot \frac{dr^i_t}{dt} = u^i_{0,t}(\rho^i_t - q_t) \cdot c^i(\rho^i_t - q_t) > 0$ unless $\rho^i_t = q_t$. Therefore $d(\sum u^i_t)/dt > 0$ unless $\rho^i_t = q_t$ for all i. This, and the continuity of the differential equation system allows us to use $\sum u^i$ as a Lyapunov function and apply the standard asymptotic

[45] There is a close correspondence between these dynamics and those found in Champsaur and Cornet (1990). Their agents also choose locally to maximize gains. However, at each point in time a *local Walrasian equilibrium* is attained.

[46] This is essentially the model in Ledyard (1974). In that paper, however, the model was ad hoc. Here we have provided a micro-foundation for it.

convergence theorems. □

B.2 LME vs ABL

To see how the ABL model differs from the LME model, consider the following. Hypotheses 1-3 imply that $q_t = \frac{\sum_i b_t^i}{I}$ in both the ABL and LME models. In both models, prices always equal the average of the bids in the market. But the two models differ in how average bids relate to the underlying utility functions. Under Hypothesis 6 of the LME model, $\frac{\sum_i b_t^i}{I} = \bar{\rho}(x_t)$. Under Hypothesis 4 of the ABL model, $\frac{\sum_i b_t^i}{I} = q_{t-\Delta} + \bar{c}\Delta A^{-1}(\bar{\rho}(x_t) - q_{t-\Delta})$. That is, in LME prices immediately change to the weighted average of the willingness to pay (at new holdings). In ABL prices adjust exponentially toward the weighted average of the willingness to pay. As such, in the ABL model, prices react more slowly to changes in allocations.

The difference between LME and ABL is even starker in the CAPM environment. For the ABL model, the price dynamic in the CAPM environment is, from (16):

$$\frac{q_t - q_{t-\Delta}}{\Delta} = A^{-1}\Omega \frac{\sum c^i a^i e^i(q_{t-\Delta}, x_t^i)}{I} \tag{OA.16}$$

For the LME model in the CAPM environment,[47]

$$\frac{q_t - q_{t-\Delta}}{\Delta} = -\frac{1}{\sum c^i}\Omega^2 \sum c^i (a^i)^2 e^i(q_{t-\Delta}, x_{t-\Delta}^i). \tag{OA.17}$$

Two striking differences with respect to dynamics emerge under ABL. First, analogous to

[47]To obtain the result, (i) take first differences of (OA.12) after lagging $(q_t - q_{t-\Delta})$, then (ii) replace $\rho^i(x^i)$ with $\mu - a^i\Omega r^i$ in order to re-express the equations in terms of $r_t^i - r_{t-\Delta}^i$ and (iii) finally use (OA.10) and the formula for CAPM excess demands, namely, $e^i(q_{t-\Delta}, x_{t-\Delta}^i) = \frac{1}{a^i}\Omega^{-1}(\mu - q_{t-\Delta}) - r_{t-\Delta}^i$.

Walrasian dynamics, price changes depend on excess demands evaluated at past allocations. Second, a minus sign features in front of the equations. This implies, among others, that a commodity's price change is *opposite to* its own (weighted) excess demand. If (weighted) excess demand is positive, the price drops. Neither prediction is upheld in the data.[48]

C Speculation

To see what happens if agents were to speculate, consider the continuous-time problem of optimally adjusting the flow of trade $z_t^i = dr_t^i$ subject to a bound on the flow size $|z_t^i|^2 \leq \gamma$ and assuming that the agent believes prices follow an Itô process. Let J denote the Hamilton-Jacobi-Bellman value function (expected utility of final consumption of the commodities, as a function of the current state, consisting of current prices and current holdings). Let J_r denote the vector of partial derivatives of J with respect to the holdings $r_{k,t}^i$, and J_0 the partial derivative of J with respect to s_t^i. It can be shown that the optimal trade flows satisfy the following equations:

$$z_{k,t}^i \sim u_{0,t}^i(\rho_k^i(x_t^i) - q_{k,t}) + (J_{r,k} - J_0 q_{k,t}) \, .$$

The first term represents local optimization: desired trade flow is proportional to the gradient of the utility function, subject to the budget constraint. The second term represents speculation. Since J denotes the expected utility of the stock (holdings) of commodities at the end of trading,

[48]It may not be immediately obvious how the results are inconsistent with the second prediction, since we transformed the regressors using Ω. The price dynamics in (OA.17) imply that the regression coefficient matrix is proportional to $-\Omega$, which has negative diagonal elements. The data reject this. Note that the *off*-diagonal elements of the coefficient matrix are nonzero. But their sign changes depending on the treatment; on average (across treatments), they equal zero. In the regressions, we did not distinguish between treatments. This was not necessary under ABL after transformation of the regressors using Ω. ABL and LME therefore make similar but not identical predictions about the off-diagonal elements of the coefficient matrix: under ABL, the true coefficients are identically zero; under LME, they are zero *on average*, across treatments. If the latter had been true, then the distribution of the corresponding z-statistics would have been affected by mixing of means, and hence, flatter than observed (compare Figures 6b and 5 [Left Panel]).

$J_{r,k}$ denotes the expected marginal utility of consumption of commodity k. If, at current prices, expected marginal utility of (eventual) consumption is proportional to the price, the second term is zero, and optimal trade flow is solely determined by local utility maximization. If marginal utility of a commodity is expected to eventually be higher than the current price (modulo J_0), the second term is positive, and the agent trades more than is needed for local maximization. Eventually, marginal utilities of consumption need to be aligned with prices: at the end of trading, i.e., at some distant T, $J_{r,k} = J_0 q_{k,T}$. If at current prices, $J_{r,k} > J_0 q_{k,t}$, the agent must expect future prices ($q_{k,T}$) to be higher than today's ($q_{k,t}$; again, we are ignoring changes in J_0). Our agent therefore speculates: she trades to a higher quantity (stock) of the commodity than she expects to eventually want; she will later reduce holdings and profit from the expected price increase.[49]

D ABL in CAPM

The price dynamics implied by our model in discrete time, see (10), are:

$$q_t - q_{t-\Delta} = \bar{c}\Delta A^{-1}\left(\mu - \Omega\frac{\sum_i a^i c^i r_t^i}{\sum_i c^i} - q_{t-\Delta}\right) \tag{OA.18}$$

Since we want to compare this to the Walrasian model (3), we write it in terms of excess demand functions. To find the Walrasian excess demand functions, maximize $\mu \cdot r^i - \frac{a^i}{2} r^i \cdot$

[49]See Sundaresan (1989) or Constantinides (1990) for analogous applications of Itô calculus to deriving optimal trade flow when utility depends on the cumulative stock (holdings).

$(\Omega r^i) - q \cdot r^i$. At lagged prices, the individual Walrasian excess demand functions are

$$e^i(q_{t-\Delta}, x_t^i) = \frac{1}{a^i}\Omega^{-1}(\mu - q_{t-\Delta}) - r_t^i. \tag{OA.19}$$

From (OA.19), $c^i a^i \Omega r_t^i = c^i(\mu - q_{t-\Delta}) - c^i a^i \Omega e^i(q_{t-\Delta}, x_t^i)$. Substituting this into (OA.18), and dividing by Δ, yields:[50]

$$\frac{q_t - q_{t-\Delta}}{\Delta} = A^{-1}\Omega \frac{\sum c^i a^i e^i(q_{t-\Delta}, x_t^i)}{I} \tag{OA.20}$$

E Newton-Raphson Algorithm vs ABL

It has often been said that the CDA is a computational device for finding the competitive equilibrium prices and allocations (Bossaerts and Plott, 2008). This is because prices in CDA experiments, without income effects and with one commodity plus numeraire, sometimes seem to mimic the Newton-Raphson (NR) algorithm which computes the zeros of a set of equations. To compute the $p^* = (1, q^*)$ that satisfies $E(q^*, \omega) = 0$ (i.e., to compute equilibrium prices), the NR algorithm does the following sequential computation:

$$q_t - q_{t-\Delta} = \left[\frac{\partial E(q_{t-\Delta}, \omega)}{\partial q}\right]^{-1} E(q_{t-\Delta}, \omega) \tag{OA.21}$$

For the CAPM model, $\frac{\partial E(q_{t-\Delta},\omega)}{\partial q} = \sum_i \frac{1}{a^i}\Omega^{-1}$. Therefore, $q_t - q_{t-\Delta} = \hat{a}\Omega \sum_i e^i(q_{t-\Delta}, \omega)$, where $\hat{a} = \left[\sum_i \left(\frac{1}{a^i}\right)\right]^{-1}$. The similarity of this to (16) is interesting. The Hessian of the utility

[50] (16) does not imply causation. That is, prices are not "responding to excess demands". It is simply a feature of the quadratic preferences of the CAPM model that let us write price changes for the ABL model this way. The theory merely says that the path of prices will satisfy (16).

functions plays a key role in both. However, the two are different. In ABL the *weighted* excess demand curves are important while in NR they are not. ABL follows a different path from NR.

F Experiments: Setup and Additional Evidence

F.1 The Structure of Market Experiments

For those unfamiliar with market experiments, a brief introduction follows. Participants are solicited, usually via email invitations, to come and participate in an experimental session at a given location (or, in some instances, access the experiment online) and at a given time. Each experimental session starts with an instructional period, where the rules of engagement are explained, participants are given the opportunity to ask questions, familiarize themselves with the trading software and participate in a practice trading session. An experiment proceeds in a series of replications, called periods. At the beginning of a period each participant i is given an initial endowment of commodities (or financial assets), w^i. Markets open and participants are free to trade subject to the usual budget constraints. Trading occurs via a market institution of the experimenter's choice. At the end of a period, participant i will have traded d^i and will have final holdings of $x^i = w^i + d^i$. Participants receive payments according to a payoff function $u^i(x^i)$, specified by the experimenter and presented to the participants during the instructional period. In some sessions all periods are payoff-relevant. In others, participants go through several periods and only some are chosen at random to be payoff-relevant.

Two standard trading institutions used in experiments are the Continuous Double Auction (CDA) and the Call Market (CM). The CDA is a trading process in which participants post limit buy and sell offers by specifying quantity and price (for example, a limit buy offer is an offer the buy a specified quantity at or below the offer price; offers are usually valid until canceled or executed, i.e., there is usually no option to have the offers lapse). In most cases the offers are displayed in an open book, i.e, they are visible to all participants. When someone submits a buy offer (bid) with a limit price above that of the best sell offer (ask) in the book, a trade takes place, at the standing offer limit price. Conversely, when someone submits a sell offer (ask) with a limit price below that of the best buy offer (bid) in the book, a trade takes place, at the bid limit price. When accepted an offer becomes part of a transaction and it is withdrawn from the order book. The CDA can be thought as an example of a system that facilitates *non-tatonnement* dynamics.

In a call market, participants also post buy and sell offers by specifying quantity and price but, contrary to the CDA, no transaction occurs or is accepted until the market is "called." If the book is closed (i,e, subjects cannot see each others' bids), this is just a sealed bid auction. If the book is open (i.e., participants can see each others' bids) and subjects can withdraw their bids and submit new ones, the call market becomes an example of a system that facilitates *tatonnement* dynamics.

In the paper we report on periods in the experiment when trade took place using the CDA. In Sessions 5–9, trade in some periods took place using a call mechanism. We do not include those periods in the analysis.

F.2 Additional Experimental Evidence

We here provide further demonstration that the cross-asset effects of excess demand on price changes replicates in four-asset experiments and even if CAPM preferences are not induced, but risk is actually realized. This means that participants come with home-grown preferences. From a pricing point of view, this does not matter: standard asset pricing results such as CAPM emerge even if uncertainty is explicit rather than induced through a nonlinear payoff function (for background literature, see Biais, e.a. (2017); Bossaerts and Plott (2004); Bossaerts, Plott and Zame (2007); Crockett, Friedman and Oprea (2017)).

The experiment was designed as follows. Three sessions were run at Caltech using the Marketscape interface (same interface as for Sessions 1–4 in the paper). There were four assets. Three of them, called A, B and C, had a random payoff depending on the drawing of one of four states. The fourth asset, the Note, was risk-free. In addition, cash was available, which was to be used in buying and selling shares in the assets. The relation between states and asset payoffs was as follows. States were equally likely to be drawn at the end of a period.[51]

State	X	Y	Z	W
A	30	190	500	200
B	100	270	300	130
C	200	210	90	180
Note	100	100	100	100

The realization of the state was unknown to participants for the duration of the period but the payoff distribution from which it was drawn (i.e., the table above) was public information. The number of participants per experiment ranged from 29 to 70.

One can readily deduce expected payoffs. They were 230, 200 and 170, for A, B and C respectively. The payoff covariance matrix can also be derived. Notice that, unlike in the example market experiment discussed in Section 3, payoff variances are *unequal*.

[51] Notional currency, called "francs," was used in all experiments. At the end of each experiment each participant's cumulative earnings from all periods were converted to US dollars at a pre-specified exchange rate ($0.04 per franc).

Ω	A	B	C
A	28850	11575	-7375
B	11575	7450	-2225
C	-7375	-2225	2250

Each session was organized as six to eight replications of the same situation. The Notes could be held in positive or negative amounts, i.e. short selling of Notes was allowed.[52] In contrast, the risky securities A, B, and C could only be held in non-negative amounts, i.e. they could not be sold short. In the beginning of each period the assets were allocated across subjects as shown in Table OA.1. Cash was allocated against a loan. This leverages their position and increases the risk of the endowments to the participants. As a result, trade takes place because of risk aversion.

Table OA.1: Experimental design data

Experiment	Participant Category Number	Signup Reward (franc)	Endowments A	B	C	Notes	Cash (franc)	Loan Repayment (franc)	Exchange Rate $/franc
991026	13	0	4	0	5	0	400	2075	0.04
	16	0	0	6	5	0	400	2350	0.04
001030	46	0	4	0	5	0	400	2075	0.04
	22	0	0	6	5	0	400	2350	0.04
001106	47	0	4	0	5	0	400	2075	0.04
	23	0	0	6	5	0	400	2350	0.04

No participant was given information regarding the asset allocations of the other participants. In each period the markets were open for a pre-set amount of time, usually ranging from 15 to 25 minutes. During open markets, the subjects had the opportunity to trade securities for cash, and thus re-balance their initial portfolios, through an online, continuous, anonymous

[52] When selling short a Note, the seller promises to pay the face value of the Note to the buyer when the Note expires. Effectively, the seller borrows the purchase price; the face value of the Note acts as a loan amount, inclusive of interest.

open-book system (Marketscape). At the end of each period each subject had his/her final portfolio of risky assets, Notes, and cash. Notice that Notes and cash were perfect substitutes in the end of a period. However, because assets could only be traded for cash, cash also had transaction value during the trading. When a period closed the state was announced and earnings recorded, to be paid out at the end of the experiment in real cash. New allocations of the assets were distributed and a fresh period opened (earnings from previous periods were NOT available as cash in subsequent periods). Subjects whose earnings were sufficiently low were declared bankrupt and were prevented from participating in subsequent periods.[53] Earnings ranged from nothing (the bankrupt participants) to over two hundred dollars.

Table OA.2: Evidence of cross-security effects in three sessions of a four-asset CAPM experiment where uncertainty was explicit rather than induced through CAPM quasi-linear payoff functions.

Exp.	Asset	Coefficients[54]				R^2	F-stat.[55]
		Constant	Excess Demand[56]				
			A	B	C		
991026	A	3.767*	1.918*	0.838*	-0.473*	0.024	5.89
($N = 710$)		(1.814)	(0.898)	(0.408)	(0.220)		
	B	1.784*	0.639	0.425*	-0.123	0.031	7.64
		(0.997)	(0.480)	(0.231)	(0.115)		
	C	-2.039**	-0.914*	-0.467**	0.214*	0.019	4.51
		(0.878)	(0.406)	(0.204)	(0.096)		
001030	A	2.556**	2.933**	1.085**	-0.775**	0.062	21.63
($N = 982$)		(0.789)	(0.921)	(0.358)	(0.240)		
	B	0.466*	0.026	0.115	0.020	0.020	6.70
		(0.249)	(0.239)	(0.091)	(0.065)		
	C	-0.336	-0.223	-0.032	0.076	0.008	2.75
		(0.763)	(0.746)	(0.300)	(0.192)		
001106	A	0.687*	0.492**	0.205**	-0.122**	0.012	6.22
($N = 1556$)		(0.416)	(0.198)	(0.091)	(0.049)		
	B	0.692*	0.174	0.168*	-0.018	0.019	10.11
		(0.370)	(0.143)	(0.083)	(0.032)		
	C	-1.031**	-0.376**	-0.152**	0.100**	0.009	4.84
		(0.282)	(0.110)	(0.051)	(0.028)		

Table OA.2 shows the results from OLS projections of price changes on excess demands after each trade. Many cross-asset slope coefficients are significant. When significant, the slope

[53]For a participant to be declared bankrupt he/she had to have negative cumulative earnings for two consecutive periods. See also Bossaerts and Plott (2004).

coefficients have the same sign as the corresponding element in the covariance matrix. E.g., the excess demand of B for instance, correlates positively with subsequent transaction price changes in A, which reflects the negative covariance between the payoffs of A and B. With one exception, the insignificant coefficients also have the right sign. This corroborates the findings in the paper for an experiment where quasi-linear preferences were not induced, but effectively obtained through uncertainty and risk aversion, and where there were four assets, not three.

G Experimental Parameters for Sessions 5-9

Session 100726	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 1 (9)	Varying Across Periods		0.06
Type 2 (9)			0.1
Securities:			
Exp Payoff ($)	0.75	0.75	
Pay Variance	1.0	0.5	
Period 1:			
Pay Covariance	-0.25		
Market (Units)	4	3	
CAPM Price	0.51	0.71	
Period 2:			
Pay Covariance	-0.25		
Market (Units)	4	3	
CAPM Price	0.51	0.71	
Period 3:			
Pay Covariance	-0.25		
Market (Units)	4	3	
CAPM Price	0.51	0.71	
Period 4:			
Pay Covariance	0.2		
Market (Units)	4	3	
CAPM Price	0.41	0.58	
Period 5:			
Pay Covariance	0.2		
Market (Units)	4	3	
CAPM Price	0.41	0.58	
Period 6			
Pay Covariance	0.2		
Market (Units)	4	3	
CAPM Price	0.41	0.58	

Session 100816	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):			
Type 1 (10)	Varying Across Periods		0.06
Type 2 (9)			0.1
Securities:			
Exp Payoff ($)	0.75	0.75	
Pay Variance	1.0	0.5	
Period 1:			
Pay Covariance	0.2		
Market (Units)	3.79	3.16	
CAPM Price	0.42	0.58	
Period 2:			
Pay Covariance	0.2		
Market (Units)	3.79	3.16	
CAPM Price	0.42	0.58	
Period 3:			
Pay Covariance	0.2		
Market (Units)	4	3	
CAPM Price	0.41	0.58	
Period 4:			
Pay Covariance	-0.25		
Market (Units)	4.21	2.84	
CAPM Price	0.49	0.72	
Period 5:			
Pay Covariance	-0.25		
Market (Units)	4.21	2.84	
CAPM Price	0.49	0.72	
Period 6			
Pay Covariance	-0.25		
Market (Units)	4	3	
CAPM Price	0.51	0.71	

Table OA.3: Parameters: Sessions 5-6. All accounting was done in U.S. dollars.

Session 101118	Securities A	Securities B	Risk Av. (a^i)	**Sessions 110608, 110609**	Securities A	Securities B	Risk Av. (a^i)
Subjects (#):				*Subjects (#):*			
Type 1 (10)	Varying Across		0.06	Type 1 (17)	Varying Across		0.06
Type 2 (8)	Periods		0.1	Type 2 (17)	Periods		0.1
Securities:				*Securities:*			
Exp Payoff ($)	0.75	0.75		Exp Payoff ($)	0.75	0.75	
Pay Variance	1.0	0.5		Pay Variance	1.0	0.5	
Period 1:				*Period 1:*			
Pay Covariance	-0.4			Pay Covariance	-0.4		
Market (Units)	4	3		Market (Units)	4	3	
CAPM Price	0.55	0.76		CAPM Price	0.54	0.76	
Period 2:				*Period 2:*			
Pay Covariance	-0.4			Pay Covariance	-0.4		
Market (Units)	4.44	2.67		Market (Units)	4	3	
CAPM Price	0.50	0.78		CAPM Price	0.54	0.76	
Period 3:				*Period 3:*			
Pay Covariance	-0.4			Pay Covariance	-0.4		
Market (Units)	4.44	2.67		Market (Units)	4	3	
CAPM Price	0.50	0.78		CAPM Price	0.54	0.76	
Period 4:				*Period 4:*			
Pay Covariance	0.4			Pay Covariance	0.4		
Market (Units)	4	3		Market (Units)	4	3	
CAPM Price	0.37	0.52		CAPM Price	0.36	0.52	
Period 5:				*Period 5:*			
Pay Covariance	0.4			Pay Covariance	0.4		
Market (Units)	3.56	3.33		Market (Units)	4	3	
CAPM Price	0.39	0.52		CAPM Price	0.36	0.52	
Period 6				*Period 6*			
Pay Covariance	0.4			Pay Covariance	0.4		
Market (Units)	3.56	3.33		Market (Units)	4	3	
CAPM Price	0.39	0.52		CAPM Price	0.36	0.52	

Table OA.4: Parameters: Sessions 7-9. All accounting was done in U.S. dollars.

H Experiments: Sample Instructions

H.1 Instructions: Session 110609 (Type 2 Subject)

Instructions

Contents:
The Experiment
The Markets Interface, Flex-E-Markets

I. The Experiment

1. Situation

The experiment consists of a number of replications of the same situation, referred to as periods. At the beginning of each period, you will be given securities and cash. Markets will open and you will be free to trade your securities. You buy securities with cash and you get cash if you sell securities. At the end of the period, the securities expire. They will pay dividends, which depend on how many securities you are holding at market close, and in which combination, as specified below.

Your period earnings has two components: the dividends on the securities you are holding after markets close, plus your cash balance.

Period earnings are cumulative across periods. There will be 12 periods in this experiment and each period lasts 5 minutes. You will be paid for twice of what you earn in 5 randomly pre-selected periods, which will be announced at the end of the experiment. The cumulative earnings are yours to keep, in addition to a standard $5 sign-up reward.

During the experiment, accounting is done in fake dollars. One fake dollar is worth 1 U.S. cent. So, 100 fake dollars is worth 1 U.S. dollar. The symbol $ is used throughout to denominate fake dollars.

2. The Securities

You will be given two types of securities, stocks and bonds. Bonds pay a fixed dividend at the end of a period, namely, $100. Stocks pay dividends that depend on the *number* of units of each you are holding and in which *combination*.

There are two stocks, A and B. At the beginning of each period, you will be given a look-up dividend table that allows you to determine the dividends that are promised for each possible combination of holdings of A and B. An example of such a look-up table is reproduced here.

8	*440*	*542*	*634*	*716*	*788*	*850*	*902*	*944*	*976*
7	*403*	*501*	*589*	*667*	*735*	*793*	*841*	*879*	*907*
6	*360*	*454*	*538*	*612*	*676*	*730*	*774*	*808*	*832*
5	*313*	*403*	*483*	*553*	*613*	*663*	*703*	*733*	*753*
4	*260*	*346*	*422*	*488*	*544*	*590*	*626*	*652*	*668*
3	*203*	*285*	*357*	*419*	*471*	*513*	*545*	*567*	*579*
2	*140*	*218*	*286*	*344*	*392*	*430*	*458*	*476*	*484*
1	*73*	*147*	*211*	*265*	*309*	*343*	*367*	*381*	*385*
0	*0*	*70*	*130*	*180*	*220*	*250*	*270*	*280*	*280*
B / **A**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**

For instance, if you are holding 2 units of A and 3 units of B at market close, the dividends on this combination of A and B will amount to $357. Or if you're holding none of A and 3 of B, the dividends will equal $203. If you're finishing with 7 units of A and 4 units of B, you'll receive $652.

Each period, the dividend table will be different. So, it is important that you pay careful attention to it before you start trading.

Although this may be of little relevance to you, you may want to know that dividend tables will generally differ across market participants.

II. The Markets Interface, *Flex-E-Markets*

You trade through an electronic market interface called *Flex-E-Markets*. Navigate to http://filagora.caltech.edu/fm/ and enter the login ID and password you have been given at the beginning of the experiment. Then go to "Marketplace Access" and pick the appropriate Marketplace. You can enter marketplace "practice" and play with various functions of Flex-E-Markets while the instruction is read.

Once you enter a marketplace, you will see slide bars for each market (Stocks A and B, and Bonds). The number of units of each security you have is displayed next to the market name. When choosing a bar, the order form will be populated. The price changes as you slide the ring on the bar. Use the order form to submit orders to buy, to sell, or to cancel previously entered orders. You can submit multiple orders at a particular price by changing the quantities in the order form. Submitted orders will show as red (if a sell order) or blue (if a buy order) tag on the slide bar. Along with your own orders, you will be able to see other participants' orders, but you will not be told who submitted those.

The orders you submit are *limit orders*. This means that, if they can be executed (i.e., if the system can find a counter party), you will be able to trade at the price you indicated, or at a better price.

How you may be able to get a better price depends on the trading mechanism. During the experiment, we will alternate across periods between two trading mechanisms: the continuous markets and the call market.

Page 3 of 4

- In the *continuous market*, *Flex-E-Markets* constantly tries to match incoming buy and sell orders. If a buy order arrives with a price at or above the highest possible price of a standing (i.e., previously entered) sell order, then the buy order trades with this best sell order, *at the price of the sell order.* Conversely, if a sell order arrives with a price at or below the highest possible price of a standing buy order, then the sell order trades with this best buy order, *at the price of the buy order.* If there are many "best" standing orders against which an arriving order can be executed, then *Flex-E-Markets* will choose the oldest standing order.
- In the *call market*, limit orders are accumulated over time without *Flex-E-Markets* trying to match. Only at the end of the period will *Flex-E-Markets* execute orders. It does so by ranking orders by price and matching high price buy orders with low price sell orders until there are no more matches for which the buy price is at least as high as the sell price. All orders execute at the sell price of the last match or the buy price of the next (unexecuted) match, whichever is higher. During order submission, flex-e-markets will periodically compute provisional clearing prices and post them. The provisional clearing prices provide an indication of the level of prices at which trade would take place if flex-e-markets were to try to clear all standing orders.

Your holding of a security is displayed above the corresponding slide bar. Your cash holding is displayed in the upper right hand corner.

If you submit an offer to buy, you need to have enough cash.

When you submit an offer to sell, you need to have enough securities.

Still, we allow you to sell bonds (but not stock) that you don't own. This is called short selling. In that case, if the sale goes through, you end up with a negative position in the bonds, and, per unit, the dividend of the bond ($100) will be *subtracted* from your total pay at the end of the period.

Because you need to have enough cash to buy, we generally start you out with lots of cash. Be careful: this cash is really "on loan," because it will be offset with a short (negative) position in the bond.

The *Flex-E-Markets* interface contains more functionality than described above (such as display of the list of orders or "books" in table format, or graphical display of incoming orders and past trades, etc.). Participants are invited to explore this functionality during the practice periods. None of it is crucial to successfully trade.

H.2 Snapshot of Online Trading Interface and Payoff Table: Session 020528 (Type 1 Subject)

MARKET SUMMARY **ID: 123 Wed Sep 10 17:08:25 2003** Period **9** -Closed- **RELOAD CURRENT DATA**

Market	Your Units	Best Buy Offer	Best Sell Offer	Last Trade	My Offers	My Trades	Graph	History
Notes	0	-@-	-@-	-	-/-			
SecurityA	2	-@-	-@-	-	-/-			
SecurityB	8	-@-	-@-	-	-/-			

Order Form
Buy Sell Market:
Units: Price:
Time to Expire: 0
(e.g. 1h6m5s; 0=never expire)
Order Clear

Your cash on hand is: **400**
Home | Instructions and Help | Inventory | Graph of All Markets | Payoff Summary | Announcements | LOGOUT

Payoff From a and b

Payoff Summary

B	Payoff																				
20	3958	4196	4433	4668	4902	5134	5365	5594	5822	6048	6273	6496	6718	6938	7157	7374	7590	7804	8017	8228	8438
19	3762	4000	4236	4471	4704	4936	5166	5395	5622	5848	6073	6295	6517	6736	6955	7172	7387	7601	7813	8024	8233
18	3566	3803	4039	4274	4506	4738	4968	5196	5423	5648	5872	6094	6315	6535	6752	6969	7184	7397	7609	7819	8028
17	3370	3607	3842	4076	4308	4539	4769	4996	5223	5448	5671	5893	6113	6332	6550	6766	6980	7193	7404	7614	7823
16	3173	3410	3645	3878	4110	4340	4569	4797	5023	5247	5470	5692	5912	6130	6347	6562	6776	6989	7200	7409	7617
15	2976	3212	3447	3680	3911	4141	4370	4597	4822	5046	5269	5490	5709	5927	6144	6359	6572	6784	6995	7204	7411
14	2779	3015	3249	3482	3713	3942	4170	4397	4622	4845	5067	5288	5507	5725	5941	6155	6368	6580	6790	6998	7205
13	2582	2817	3051	3283	3514	3743	3970	4196	4421	4644	4866	5086	5304	5522	5737	5951	6164	6375	6585	6793	6999
12	2385	2620	2853	3084	3314	3543	3770	3996	4220	4443	4664	4884	5102	5318	5533	5747	5959	6170	6379	6587	6793
11	2187	2421	2654	2885	3115	3343	3570	3795	4019	4241	4462	4681	4899	5115	5330	5543	5754	5965	6173	6381	6586
10	1990	2223	2456	2686	2916	3143	3370	3594	3818	4039	4260	4478	4696	4911	5126	5338	5550	5759	5968	6174	6380
9	1791	2025	2257	2487	2716	2943	3169	3393	3616	3837	4057	4275	4492	4707	4921	5133	5344	5554	5761	5968	6172
8	1593	1826	2057	2287	2516	2743	2968	3192	3414	3635	3854	4072	4288	4503	4717	4929	5139	5348	5555	5761	5965
7	1395	1627	1858	2088	2315	2542	2767	2990	3212	3432	3651	3869	4085	4299	4512	4723	4933	5142	5349	5554	5758
6	1196	1428	1659	1888	2115	2341	2565	2788	3010	3230	3448	3665	3881	4095	4307	4518	4727	4935	5142	5347	5550
5	997	1229	1459	1687	1914	2140	2364	2586	2807	3027	3245	3461	3676	3890	4102	4312	4521	4729	4935	5139	5342
4	798	1029	1259	1487	1714	1939	2162	2384	2605	2824	3041	3257	3472	3685	3897	4107	4315	4522	4728	4932	5134
3	599	830	1059	1286	1512	1737	1960	2182	2402	2620	2838	3053	3267	3480	3691	3901	4109	4315	4520	4724	4926