WORKING PAPERS

N° 798

August 2024

“Asset Pricing and Risk Sharing in CompleteAsset Pricing and Risk Sharing in Complete Markets: An Experimental Investigation”

Bruno Biais, Thomas Mariotti, Sophie Moinas, Sébastien Pouget

Toulouse School of Economics

Asset Pricing and Risk Sharing in Complete Markets: An Experimental Investigation*

Bruno Biais† Thomas Mariotti‡ Sophie Moinas§ Sébastien Pouget¶

August 17, 2024

Abstract

We study asset pricing and risk sharing in experimental financial markets. We design our experiment to test the key equilibrium implications of rational choice and competitive behavior in complete markets without making parametric assumptions on preferences. We find that participants behave competitively but deviate from rationality, as around 25% of their actions are first-order stochastically dominated. We propose a random-choice model predicting that, as the number of participants grows large, prices and average per-participant trades converge to those in the rational-choice competitive equilibrium. This prediction is supported by our experimental data. We structurally estimate a special case of the random-choice model with CRRA utilities and logit weighting functions and find that only around 80% of participants benefit from participating in the market.

Keywords: Asset Pricing, Risk Sharing, Experimental Financial Markets, Complete Markets, Convergence to Equilibrium, Random-Choice Model.
JEL codes: G12, C92.

*We thank Elena Asparouhova, Peter Bossaerts, George Constantinides, John Duffy, Frank Heinemann, Debrah Meloso, Charles Plott, Martin Weber, Bill Zame, and Stefan Zeisberger for extremely valuable feedback. This research has benefited from financial support of the ANR (Programme d'Investissements d'Avenir ANR-17-EURE-0010).

†HEC Paris. Email: `biaisb@hec.fr`.

‡Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, CEPR, and CESifo. Email: `thomas.mariotti@tse-fr.eu`.

§University of Toulouse Capitole, Toulouse School of Management, and Toulouse School of Economics. Email: `sophie.moinas@tse-fr.eu`

¶University of Toulouse Capitole, Toulouse School of Management, and Toulouse School of Economics. Email: `sebastien.pouget@tse-fr.eu`

1 Introduction

Since the seminal works of Debreu (1959) and Arrow (1964), the theory of competitive equilibrium in complete markets with rational agents has offered an elegant framework with sharp implications: agents should share risk perfectly and only aggregate risk should be priced (Borch (1962)). This theory led to a cornerstone of financial economics: the consumption-based capital asset-pricing model (Rubinstein (1976), Lucas (1978), Breeden (1979)). However, many empirical studies based on field data rejected this model (see, e.g., the surveys by Campbell (2003) and Ludvigson (2013)).

Is the theory rejected because real-world asset markets are imperfect and incomplete (Grossman and Stiglitz (1980), Geanakoplos and Polemarchakis (1986))? Is it rejected because it is difficult to measure risk in the field (Rietz (1988), Julliard and Ghosh (2012))? Or is it rejected because human cognition and preferences do not conform to the standard rational-choice paradigm (DeBondt and Thaler (1985), Shiller (2000))? It is important to tell which of these three explanations is at play, because they have very different implications for the study of financial markets: the first one calls for better models of market imperfections, the second one for better empirical measures of risk, and the third one for alternative models of human decision making. However, disentangling these three explanations using field data is difficult because of measurement and identification issues.

We investigate these explanations using laboratory experiments in which markets are complete and the level of individual and aggregate risk is controlled. In that context, any deviation from rational-choice competitive equilibrium can be attributed to either noncompetitive or imperfectly rational behavior. We design our experiments so that the two hypotheses that participants are competitive and that they are rational can be tested without making parametric assumptions on preferences. In our simple setting, with two equiprobable states of nature and complete markets, without specifying preferences we can identify if actions are first-order stochastically dominated, and also if they are second-order stochastically dominated.

Our experimental market is designed to emulate the competitive-equilibrium framework: we elicit supply and demand schedules from participants, which we can then aggregate, and cross to determine market-clearing prices. We implement two within-subject experimental treatments. The first treatment dimension is related to risk: there is aggregate risk in some replications but not in others. This enables us to test the theoretical prediction that a risk premium should arise if and only if there is aggregate risk. The second treatment dimension is related to the market mechanism: the price is set by market clearing in some replications whereas it is drawn at random in others, in the spirit of Becker et al. (1964). Comparing the market-clearing and random-price treatments enables us to test, and reject, the hypothesis that participants behave noncompetitively.

We ran the experiment with 220 students at Toulouse University. There were thirteen cohorts. Each cohort participated in a session consisting of eight independent replications of the experimental market. Participants' compensation was a linear function of their gains in two randomly drawn replications. Payment per participant varied from €5 to €120, with an average of €85.84. This is a significant amount relative to the average participant's monthly budget of €646.

At the individual level, we observe deviations from rationality: Around 25% of the participants' actions are first-order stochastically dominated.[1] Such irrational actions lead to noisy prices in small markets. Yet, simulations based on our experimental data show that, as the number of participants grows large, market-clearing prices converge to equilibrium: when there is no aggregate risk, the market-clearing price of the stock converges to its expected dividend, whereas when there is aggregate risk there is a risk premium.

To reconcile these market- and individual-level findings, we develop and test a random-choice model of bounded rationality in the spirit of Luce (1959) and McKelvey and Palfrey (1995, 1998). In this model, the probability that an agent selects an action is increasing in her utility from that action.[2] Our random-choice model is nonparametric in that it does not rely on parametric specifications of the agents' preferences: the only assumptions are that agents' utility is lower when they take first-order or second-order stochastically dominated actions than when they take nondominated actions. Our random-choice model predicts that, as the number of participants grows large, individual deviations from rationality should average out and aggregate outcomes should converge to equilibrium. This is in line with our experimental findings. Our random-choice model also predicts that dominated actions ("mistakes") should be less frequent than nondominated ones, and that large mistakes should be less frequent than small ones. Again, this is in line with our experimental findings.

In order to quantify allocative efficiency in our experimental market, we then structurally estimate the special case of our random-choice model arising for CRRA utility and logit weighting. We find that around 80% of participants are better off participating in the market than staying in autarky, whereas the remaining 20% are worse off. The latter are frequently buyers purchasing too many shares of the stock at too high prices.

Related Literature Our work is in line with the seminal analyses of Bossaerts and Plott (2004) and Bossaerts et al. (2007),[3] with four key differences.

[1]Dominated actions, however, become less frequent with experience, a sign of learning.

[2]It is in general difficult to bring this type of model to the data, especially in the context of continuous double auctions. In our simple experimental setting, however, agents actions are just the quantity offered or demanded at any given price, which lends itself to random-choice modeling.

[3]Whereas our theoretical framework is static, Bossaerts et al. (2015), Asparouhova et al. (2016), and Crockett et al. (2019) offer interesting analyses of equilibrium market dynamics.

First, whereas Bossaerts and Plott (2004) and Bossaerts et al. (2007) study a continuous double auction, we strive to emulate as closely as possible the competitive-equilibrium framework. Eliciting supply and demand schedules and crossing them to determine market-clearing prices enables us to test the theory, estimate model parameters, run simulations, and compute counterfactuals.

Second, whereas Bossaerts and Plott (2004) and Bossaerts et al. (2007) consider a three-state three-asset environment, in which they study diversification, we only consider a two-state two-asset environment. This enables us to to avoid the difficulties associated with the design of experimental markets in which several assets are traded simultaneously (Bossaerts et al. (2002) and Asparouhova et al. (2024).) Moreover, focusing on the case of equiprobable states enables us obtain sharp theoretical implications without making parametric assumptions on participants' preferences.

Third, whereas Bossaerts and Plott (2004) and Bossaerts et al. (2007) consider markets in which there is always aggregate risk, we also consider a treatment in which there is no aggregate risk but individual endowments are risky. This enables us to test the hypothesis that, in complete markets with no aggregate risk, agents should bear no risk in equilibrium, which should give rise to risk-neutral pricing.

Fourth, whereas Bossaerts et al. (2007) model an individual trader's demand function as the sum of a mean-variance optimal demand and of a noise term to capture unobserved heterogeneity in preferences, we use a random-choice model to capture bounded rationality. The random-choice model allows for a large class of preferences, such as expected utility with arbitrary curvature or rank-dependent expected utility. In our simple experimental design, the random-choice model imposes structure on the distribution of individual decisions, which we confront to the experimental data.

Our analysis is also in line with the study of Crockett et al. (2021), who show that aggregation pathologies in financial markets become less frequent as participants' preferences become more heterogeneous. Instead of manipulating preference heterogeneity as in Crockett et al. (2021), we vary the number of participants. This enables us to show that aggregate outcomes converge to equilibrium as the number of participants grows large.

Fattinger (2021) uses our experimental framework to study how ambiguity affects asset pricing. He extends our design to include an asset with an ambiguous dividend and studies how ambiguity affects supply and demand schedules, as well as the price of the asset.

The present paper is also related to the experimental literature studying the consequences of risk aversion for economic and financial decisions. Holt and Laury (2002) observe lottery choices consistent with risk aversion. Experimental findings also suggest that risk aversion is prevalent in private-value auctions (Goeree et al. (2002)), asymmetric matching-penny games

(Goeree et al. (2003)), and one-shot matrix games (Goeree and Holt (2004)). The majority of participants in Holt and Laury (2002) have estimated relative risk aversion between 0.15 and 0.97. Estimates of the same order of magnitude have been found by Goeree et al. (2002, 2003) and Goeree and Holt (2004). Our estimates are consistent with these findings, although they are obtained in a very different economic environment.

Several interesting experimental studies focus on investment-allocation decisions (see, e.g., Kroll et al. (1988), Kroll and Levy (1992), and Magnani et al. (2022)). In contrast with these analyses, our paper studies how investment decisions determine market-clearing prices, as well as the pricing and allocative efficiency of market outcomes.

Finally, our finding that participation in financial markets can be detrimental to welfare is related to the findings of Barber and Odean (2000, 2001) that overconfidence can lead to damageable excessive trading (see also Biais et al. (2005)).

The paper is organized as follows. Section 2 presents our experimental design. Section 3 derives the testable implications of rational-choice competitive equilibrium. Section 4 offers a first analysis of our experimental data. Section 5 develops a random-choice model and confronts its implications with our experimental data. Section 6 concludes. Proofs are in the Appendix. The instructions for the experiment, as well as a robustness check, are in the Online Appendix.

2 Experimental Design

We design our experimental financial market so as to study asset pricing, risk sharing, and allocative efficiency in a setting in which the predictions from rational-choice competitive equilibrium can be tested without imposing parametric restrictions on preferences.

Assets There are two equiprobable states of nature, $\omega = u, d$, and two nonredundant assets, a stock and a bond. Hence markets are complete. We restrict our attention to equiprobable states so that theoretical predictions can be derived with minimal assumptions on preferences.[4] One unit of the bond pays 1 in each state of nature. One share of the stock pays 120 in state u and 0 in state d. To implement complete markets in our two-state experimental setting, it is enough to consider a single market on which the stock is traded for the bond; i.e., we take the bond as the numéraire.[5] We denote by S the price of the stock in units of bond.

[4] As discussed in Section 3, our predictions are based on first-order stochastic dominance, representing the choices of agents who prefer more to less consumption, and second-order stochastic dominance, representing the choices of risk-averse agents.

[5] The mapping between our experimental setting and the underlying complete-market environment is made explicit in Appendix A.1.

Endowments Participants' endowments come from initial holdings of the stock and of the bond, and from additional state-contingent income. Participants can receive three types of endowments:

1. Type 1 participants initially receive 5 shares of the stock and no bond. Their additional income is 0 in state u and 360 in state d. Thus their endowments are 600 in state u and 360 in state d.

2. Type 2 participants initially receive no share of the stock and 310 bonds. Their additional income is 0 in state u and 240 in state d. Thus their endowments are 310 in state u and 550 in state d.

3. Type 3 participants initially receive no share of the stock and 310 bonds. They receive no additional income. Thus their endowment is 310 in both states u and d.

We consider two treatments that enable us to test sharp predictions of the theory, as explained in Section 3. In Treatment I, if the number of participants is even, there are only Type 1 and Type 2 participants, in equal numbers; whereas, if the number of participants is odd, there is one additional Type 3 participant. This treatment corresponds to a situation with no aggregate risk: the sum of individual endowments is the same in state u as in state d. Type 1 and Type 2 participants can fully hedge, by trading 2 shares each.[6]

In Treatment II, if the number of participants is even, there are only Type 1 and Type 3 participants, in equal numbers; whereas, if the number of participants is odd, there is one additional Type 3 participant. This treatment corresponds to a situation with aggregate risk: the sum of individual endowments is larger in state u than in state d.

Participants do not know that there are two treatments. In any replication of the experiment, they are only informed of their own endowments and of the distribution of the dividend. As emphasized by Bossaerts and Plott (2004), this is in line with the standard competitive-equilibrium model, in which agents only rely on prices and rational expectations about the distribution of asset payoffs.

Supply, Demand, and Prices We design our experimental financial market to closely emulate the competitive-equilibrium model, in which each agent states how much she is willing to sell or buy at each price, and the price is set to clear the market.

To simplify the task of participants in the experiment, we restrict Type 1 participants to supply, and Type 2 and 3 participants to demand, nonnegative quantities of the stock. This does not affect the equilibrium when all agents are competitive and risk-averse. A further

[6] If Type 1 participants sell 2 shares at price S, then their payoff in state u is $(5-2)\times 120+2\times S$, which is equal to their payoff in state d, $(5-2)\times 0+2\times S+360$. If Type 2 participants buy 2 shares at price S, then their payoff in state u is $2\times 120-2\times S+310$, which is equal to their payoff in state d, $2\times 0-2\times S+550$.

simplification of the experiment is that participants are restricted to trade quantities no larger than 4. These restrictions forbid short selling and borrowing, which, anyhow, do not arise in equilibrium when agents are risk-averse.

Participants are asked which quantity of the stock they are willing to sell or to buy at every point of a price grid. In our baseline experiment, the price grid ranges from 52 to 62 with unit increments. In our robustness experiment, the price grid ranges from 20 to 70 with increments of 5 units.[7] These ranges include the expected value of the dividend, 60, and allow for a discount to arise in compensation for risk. Whereas the price grid is discrete, we allow participants to supply or demand any quantity of the stock in the interval $[0, 4]$. Thus participants can fine-tune their supply or demand, which enables them to equate to the price their marginal willingness to sell or to buy.

Once supply and demand schedules have been elicited, it would seem natural to simply aggregate and cross them to determine the market-clearing price. Two difficulties arise, however. First, the discreteness of the price grid may prevent the market from clearing. Second, participants may behave strategically and try to manipulate the price. To address these issues, we consider the two following mechanisms:

1. In the call-market mechanism, the price is set to minimize the gap between supply and demand.[8] This is in line with the call-market mechanism in Smith et al. (1982), McCabe et al. (1992), or Plott and Pogorelskiy (2017). The main difference is that, in these studies, participants submit schedules of limit orders, whereas in the present one they submit supply or demand schedules. Unlike those generated by limit-order schedules, the supply and demand schedules we elicit need not be monotone.[9] As mentioned above, because we consider a discrete price grid, at the price minimizing the gap between supply and demand, there is typically rationing. To make sure that this consideration does not affect the participants' choices, the experimenter supplies or demands the quantity of stock needed to clear the market. Thus participants' orders are fully executed. In practice, the additional supply or demand injected in the market by the experimenter averaged to 6.52% of the quantity traded during the call sessions. This way of handling potential mismatch between supply and demand is most conducive to participants behaving as predicted by the competitive model.[10]

[7] We design the robustness experiment to study whether our results still hold when the stakes are larger than in the baseline experiment. Indeed, because trading can occur at prices farther away from the expected value of the dividend, the cost of mistakes is larger in the robustness than in the baseline experiment.

[8] In the very few instances in which several prices minimized the gap between supply and demand, we drew one of them at random.

[9] In the standard competitive-equilibrium model too, demand functions may be nonmonotone, when wealth effects are stronger than substitution effects.

[10] Using instead a rationing rule may have induced strategic behavior. For example, participants facing prorata rationing could be tempted to inflate their demand.

2. In the random-price mechanism, the price is randomly drawn from the price grid, all prices being equiprobable. As in the call mechanism, participants' orders are fully executed. This mechanism is in line with the Becker et al. (1964) mechanism, the difference being that participants in our random-price mechanism can trade any quantity of the stock in $[0, 4]$.

At the beginning of each replication of the experiment, all participants are told which mechanism will be used to set the price. Strategic considerations could affect participants' behavior when they know that the price will be set to minimize the gap between supply and demand, but not when they know that the price will be randomly set. Therefore, comparing actions in the two mechanisms enables us to test whether participants are competitive or noncompetitive.

Implementation Our experiment includes 220 subjects, 141 for the baseline and 79 for the robustness sessions. All subjects were students enrolled in the first year of Toulouse University's master in finance.[11] There were eight cohorts in the baseline and five in the robustness experiment. All cohorts participated in eight replications of the experiment, lasting overall one and a half hour.[12] Treatment I with no aggregate risk and Treatment II with aggregate risk alternated during the eight replications. Also, for half of the cohorts, prices were randomly set in the first four replications, whereas prices were set to minimize the gap between supply and demand in the last four replications. For the other cohorts it was the other way round, with random pricing in the last four replications. The details of the experimental sessions are documented in Table 1 for the baseline sessions and in Table OA.1 in the Online Appendix for the robustness sessions.

In each cohort, two of the eight replications were randomly drawn at the end of the experiment, one from the replications in which the price was set to minimize the gap between supply and demand and one from those in which the price was randomly set. As announced at the start of the experiment, participants received the sum of their final earnings in these two replications, divided by ten. The average individual payment was €85.84, the minimum was €5, and the maximum €120. For comparison, we ran an anonymous survey among the participants, asking them their monthly budget (including all expenses: housing, food, leisure, ...). The average was €646. Thus the amount participants could make in the lab and its variability were significant relative to their budget in the field.[13]

[11]These students came from different backgrounds. The majority studied management and had very little exposure to microeconomics or financial theory. Some studied economics and had greater exposure to microeconomics, but not to financial theory. Others studied engineering or mathematics.

[12]For one cohort, an operational problem prevented us from collecting the data in the last replication.

[13]Bossaerts and Plott (2004) compare participants' behavior in experimental financial markets in the US and in Bulgaria. The experimental design is the same in the two countries but monetary incentives are much stronger in Bulgaria. They find qualitatively similar results in the US and in Bulgaria. In both countries

In the baseline experiment, 68 participants were Type 1 sellers during all replications, 68 participants alternated between being Type 2 or Type 3 buyers, and 5 participants were Type 3 buyers during all replications. In the robustness experiment, 38 participants were Type 1 sellers, 38 participants alternated between being Type 2 or Type 3 buyers, and 3 were Type 3 buyers during all replications. The instructions for our experiment are provided in the Online Appendix. The experiment was programmed on z-Tree (Fischbacher (2007)).

3 Testable Implications of Competitive Equilibrium under Rational Choice

3.1 Individual Behavior

If lottery A is first-order stochastically dominated by lottery B, an agent whose preferences are strictly increasing in state-contingent wealth prefers B to A. The corresponding class of preferences is very large. It includes expected utility or rank-dependent expected utility (Quiggin (1982)) with the only requirement that utility increases in wealth, or Fréchet differentiable utility with increasing local utility functions (Machina (1982)). We shall refer to agents with preferences in this class as FOSD agents. As explained below, our experimental setting generates clear predictions for such agents. This enables us to experimentally test the predictions from theory under mild restrictions on participants' preferences.

We also examine the implications of agents' preferences towards risk. An agent is risk-averse if she prefers any lottery to a mean-preserving spread of that lottery. Conversely, an agent is risk-loving if she prefers a mean-preserving spread of any lottery to that lottery. In addition to risk aversion, for our equilibrium analysis we also rely on second-order risk aversion (Segal and Spivak (1990)). For a small bet, the premium required by an agent exhibiting second-order risk aversion is proportional to the square of the size of that bet. Intuitively, this means that the agent is almost risk-neutral for small bets.

In our simple setting with two equiprobable states, the lotteries faced by the agents can be expressed in terms of means and standard deviations, and this formulation facilitates the comparison of the lotteries in terms of first-order stochastic dominance. Let $W^T(\omega, q, S)$ be the final wealth in state ω of a Type T agent trading a quantity q at price S. Because the two states are equiprobable, her expected final wealth is

$$\mu^T(q, S) \equiv \frac{1}{2}\,[W^T(u, q, S) + W^T(d, q, S)], \tag{1}$$

with standard deviation

$$\sigma^T(q, S) \equiv \frac{1}{2}\,|W^T(u, q, S) - W^T(d, q, S)|. \tag{2}$$

there is a risk premium. The main difference is that the risk premium is larger in Bulgaria.

So, when $W^T(u,q,S) > W^T(d,q,S)$,

$$W^T(u,q,S) = \mu^T(q,S) + \sigma^T(q,S) \text{ and } W^T(d,S) = \mu^T(q,S) - \sigma^T(q,S),$$

and, when $W^T(u,q,S) < W^T(d,q,S)$,

$$W^T(u,q,S) = \mu^T(q,S) - \sigma^T(q,S) \text{ and } W^T(d,q,S) = \mu^T(q,S) + \sigma^T(q,S).$$

Thus the lottery faced by a Type T agent trading a quantity q at price S is

$$L^T(q,S) \equiv \left(\mu^T(q,S) + \sigma^T(q,S), \frac{1}{2}; \mu^T(q,S) - \sigma^T(q,S), \frac{1}{2}\right). \tag{3}$$

In turn, this implies that, if $\mu^T(q,S) > \mu^T(q',S)$ but $\sigma^T(q,S) = \sigma^T(q',S)$, then $L^T(q,S)$ first-order stochastically dominates $L^T(q',S)$.[14]

Type 1 For a Type 1 agent, selling $2+\chi$ shares and selling $2-\chi$ shares lead to equally volatile final wealth for all $\chi \in (0,2]$. However, when the price S of the stock is strictly lower than its expected dividend of 60, selling $2+\chi$ shares leads to a lower expected final wealth than selling $2-\chi$ shares. The reverse holds when S is strictly higher than 60. By (3), this implies that supplying $2+\chi$ shares is first-order stochastically dominated by supplying $2-\chi$ shares for a Type 1 agent when $S < 60$, whereas the reverse holds when $S > 60$. Moreover, if a Type 1 agent is risk-loving or risk-neutral, it is optimal for her to supply 4 shares at any price $S > 60$ and 0 share at any price $S < 60$. Indeed, these extreme trades are those that induce the most volatile final wealth and the highest expected final wealth.

What if the price S of the stock is equal to its expected dividend of 60? At this price, a Type 1 agent is perfectly hedged if she exactly sells 2 shares. Any other trade generates volatile final wealth and the same expected wealth. Thus, a risk-averse Type 1 agent finds it optimal to exactly supply 2 shares at $S = 60$. In contrast, a risk-loving Type 1 agent prefers to take as much risk as possible at $S = 60$, which can be achieved, indifferently, by supplying 0 or 4 shares. Finally, a risk-neutral Type 1 agent is indifferent between all trades in $[0,4]$ at $S = 60$. Together, these considerations lead to our first implication, whose proof is in Appendix A.2.

Implication 1 *The following holds*:

(i) An FOSD Type 1 agent does not supply more than 2 *shares at* $S < 60$, *nor supplies less than* 2 *shares at* $S > 60$. *Moreover, if she is risk-loving or risk-neutral, then she supplies* 0 *share at* $S < 60$ *and* 4 *shares at* $S > 60$.

[14] First-order stochastic dominance obtains because the states u and d are equiprobable. For arbitrary probabilities, first-order stochastic dominance does not always obtain. Hence, equiprobability is crucial to obtain sharp testable implications. This is the reason why we use this feature in our experimental design.

(ii) A risk-averse (respectively, risk-loving) Type 1 agent supplies 2 *shares (respectively,* 0 *or* 4 *shares) at* $S = 60$*. A risk-neutral Type 1 agent is indifferent between all trades in* $[0, 4]$ *at* $S = 60$*.*

To illustrate this implication, Panel A of Figure 1 depicts possible supply functions of Type 1 agents in the price-quantity plane. The relevant quantities q are between 0 and 4, and, in the baseline experiment, the relevant prices S are between 52 and 62. There are four quadrants determined by the horizontal line $q = 2$ and the vertical line $S = 60$. It follows from Implication 1 that the North-West and South-East quadrants are dominated for an FOSD Type 1 agent. Moreover, a risk-averse Type 1 agent's supply function must go through the point $(60, 2)$.

Type 2 Following the logic we used for Type 1 agents, we obtain a symmetric implication for Type 2 agents.

Implication 2 *The following holds:*

(i) An FOSD Type 2 agent does not demand less than 2 *shares at* $S < 60$*, nor demands more than* 2 *shares at* $S > 60$*. Moreover, if she is risk-loving or risk-neutral, then she demands* 4 *shares at* $S < 60$ *and* 0 *share at* $S > 60$*.*

(ii) A risk-averse (respectively, risk-loving) Type 2 agent demands 2 *shares (respectively,* 0 *or* 4 *shares) at* $S = 60$*. A risk-neutral Type 2 agent is indifferent between all trades in* $[0, 4]$ *at* $S = 60$*.*

This implication is illustrated in Panel B of Figure 1, which mirrors Panel A.

Type 3 In contrast with Type 1 and Type 2 agents, a Type 3 agent starts with a risk-free initial endowment. As a result, trading increases her exposure to risk. Therefore, no prediction arises from first-order stochastic dominance arguments only. Nevertheless, as shown in Appendix A.2, we obtain the following weaker implication, illustrated in Panel C of Figure 1.

Implication 3 *A risk-averse Type 3 agent demands* 0 *share at* $S \geq 60$*, and a risk-loving Type 3 agent demands* 4 *shares at* $S < 60$*.*

3.2 Market Outcomes

Equilibrium outcomes at the market level depend on the treatment, i.e., on whether there is no aggregate risk, as in Treatment I, or there is aggregate risk, as in Treatment II. If all participants are FOSD second-order risk-averse agents, clear theoretical predictions obtain.

Treatment I In Treatment I, there are N Type 1 participants, N Type 2 participants, and possibly one Type 3 participant. By Implications 1–3, if all participants are risk-averse, at $S = 60$ every Type 1 participant supplies 2 shares, every Type 2 participant demands 2 shares, and every Type 3 participant demands 0 share. As a result, $S = 60$ is an equilibrium price, and the corresponding trading volume is $2N$.

Is it the only equilibrium? At price $S > 60$, by Implication 1, aggregate supply by FOSD Type 1 participants is at least $2N$. Thus, by Implications 2–3, $S > 60$ can be an equilibrium price only if at that price all Type 2 participants demand 2 shares. Now, if Type 2 participants are second-order risk-averse, then, for every small enough $\chi > 0$, buying 2 shares at $S > 60$ is dominated by buying $2 - \chi$ shares (Segal and Spivak (1990)). Indeed, relative to the safe lottery obtained when buying 2 shares, buying $2 - \chi$ shares adds a small gamble with positive expected payoff, which second-order risk-averse Type 2 participants are willing to accept. Consequently, there cannot exist an equilibrium with a price strictly above 60, and a symmetric argument rules out an equilibrium with a price strictly below 60. These remarks yield our fourth implication, proved in Appendix A.2.

Implication 4 *In Treatment I, if all participants are FOSD second-order risk-averse agents, the equilibrium is unique and such that* $S = 60$, *Type 1 participants sell* 2 *shares, Type 2 participants buy* 2 *shares, and the Type 3 participant, if any, does not trade.*

In Treatment I, aggregate wealth is constant across states. Thus, in equilibrium, risk-averse agents perfectly hedge their risk exposure and the price of the stock is simply equal to its expected dividend, so that there is no risk premium.[15] When participants are expected-utility maximizers with strictly concave differentiable utility functions, this is an instance of the mutuality principle (Borch (1962)). In our simple experimental design, it obtains as soon as participants exhibit second-order risk aversion.

Treatment II In Treatment II, there are only Type 1 and Type 3 participants. Now, by Implication 1, at any price $S \geq 60$, a risk-averse FOSD Type 1 participant supplies at least 2 shares, and, by Implication 3, a risk-averse Type 3 participant demands 0 share. Thus there exists no equilibrium such that $S \geq 60$. Nevertheless, if all participants are FOSD risk-averse agents, then their preferences are monotone and convex in state-contingent wealth, and standard equilibrium-existence results apply.[16] As shown in Appendix A.2, this

[15]Whereas the equilibrium analysis underlying Implication 4 is conducted in a perfect market, with continuous prices, it also applies in our experimental setting with a discrete price grid, because the equilibrium price $S = 60$ belongs to the price grid.

[16]Uniqueness of equilibrium is not guaranteed in general, but if participants are risk-averse expected utility maximizers with relative risk aversion at most equal to 1, then substitution effects dominate income effects and the equilibrium is unique. This is because this condition ensures that the aggregate excess-demand function for state-contingent wealth satisfies the gross-substitute property (Mas-Colell et al. (1995, Example 17.F.2 and Proposition 17.F.3)).

yields the following implication.

Implication 5 *In Treatment II, if all participants are FOSD risk-averse agents, then there exists an equilibrium such that* $S < 60$ *and the quantity traded per agent is lower than 2.*

The gap between the stock's expected dividend of 60 and its equilibrium price S is the risk premium requested by risk-averse agents to bear aggregate risk.

4 Experimental Findings

This section describes supply, demand, and prices in our experiment, tests for competitive behavior and rationality, and studies the impact of market size and learning. We focus on our baseline experiment; results from our robustness experiment are similar and are presented in the Online Appendix.

4.1 Descriptive Results

4.1.1 Aggregate Outcomes

To shed light on aggregate supply and demand in our experimental financial markets, we aggregate the individual supply and demand schedules from all Type 1 and Type 2 participants in Treatment I, and from all Type 1 and Type 3 participants in Treatment II. It is legitimate to aggregate across cohorts because, in a competitive setting, individual demand or supply should be independent of the number and characteristics of other market participants. We show below that the hypothesis that participants behave competitively cannot be rejected.

Panel A of Figure 2 depicts aggregate demand and supply in Treatment I, divided by the number of participants and replications to facilitate the interpretation. Demand is approximately decreasing and supply is approximately increasing. The price minimizing the gap between supply and demand is 60, which is the equilibrium price predicted by theory for competitive risk-averse agents. Moreover, at that price, average trading volume per participant is very close to 2 shares, which is again in line with the predictions of the theory for competitive risk-averse agents (see Implication 4).

Panel B of Figure 2 depicts aggregate demand and supply in Treatment II. Supply is not very different from its counterpart in Panel A. This is expected, as Type 1 participants face the same environment in both treatments and do not know which treatment is applied. In contrast, as predicted by theory for risk-averse agents, demand at price $S \geq 60$ is lower than in Panel A—although it does not drop to zero, which is consistent with some Type 3 participants being risk-loving or making mistakes. As a result, the price minimizing the gap between supply and demand is lower than the expected dividend of 60 and the average trading volume per participant is less than 2 shares. These results are in line with the predictions of the theory for competitive risk-averse agents (see Implication 5).

The results for the robustness experiment are offered in Figure OA.1 in the Online Appendix and are consistent with the baseline experiment.

4.1.2 Individual Decisions

To describe individual behavior, we compute, for each participant, the average supply or demand schedules across the various replications. Figure 3 plots these individual schedules depending on the type of participant and on the treatment. We observe a large heterogeneity across participants. In the following, we study if this heterogeneity can be attributed to differences in preferences among rational competitive agents, or to differences in strategic behavior, or to deviations from rationality.

4.2 Test of Strategic Behavior

To test for strategic behavior, we compare behavior when the price is set to minimize the gap between supply and demand and behavior when the price is set randomly. In the latter, by construction, agents' actions cannot affect the price so there is no scope for strategic behavior. Figure 4 depicts aggregate supply and demand in the two mechanisms, in Treatment I and Treatment II. Aggregate supply and demand schedules are similar in the two mechanisms and the figure does not suggest strategic over- or under-bidding.

To formally establish the absence of strategic behavior in our data, we propose a test based on individual behavior. We regress the quantity supplied or demanded by an individual on the indicator variable taking the value 1 when the price is set randomly and 0 otherwise. We run a panel regression with individual and period fixed effects and standard errors clustered at the individual level. We control for the proposed price and for the treatment (aggregate versus no aggregate risk). The results are in Table 2.

Table 2 shows that the dummy for random pricing is insignificant. This indicates that individual behavior does not vary with the mechanism. We can thus conclude that strategic considerations are absent from our experiment, so that our experimental market emulates a competitive market.

Table 2 also shows that individual demand schedules shift downward when there is aggregate risk, consistently with competitive rational choice theory. Indeed, as illustrated in Panels B–C of Figure 1, FOSD risk-averse buyers tend to bid lower quantities when participating as Type 3 agents, i.e., in Treatment II with aggregate risk, than as Type 2 agents, i.e., in Treatment I without aggregate risk. This result is present both without time fixed effects (the coefficient of the aggregate-risk dummy is significantly negative) and with time fixed effects.[17] As a placebo test, we focus on sellers' supply schedules. By design, sellers' asset endowment and additional incomes are constant across treatments. Rational choice theory thus predicts that their propensity to trade should be the same across treatments. Table 2 is consistent with this prediction: sellers' supply schedules are not affected by the absence or presence of aggregate risk.

4.3 Test of Imperfect Rationality

When there is no aggregate risk, Implications 1–2 offer sharp predictions for rational supply and demand behavior at the individual level: supplying more than 2 shares at $S < 60$, or less than 2 shares at $S > 60$, is first-order stochastically dominated for a Type 1 agent; symmetrically, demanding less than 2 shares at $S < 60$, or more than 2 shares at $S > 60$, is first-order stochastically dominated for a Type 2 agent.

To test whether individual behavior conforms with first-order stochastic dominance, we compute the frequency of dominated and nondominated actions at prices other than 60 for Type 1 and 2 participants.[18] As mentioned above, 68 participants in our experimental financial markets were Type 1 sellers, and each participated in eight replications of the market. For each of these participants and for each replication, we compute the proportion of dominated actions across the ten relevant prices in the price grid, i.e., across all prices from 52 to 62 except 60. We then compute, for each replication, the average of this proportion across the 68 participants. The solid line in Panel A of Figure 5 plots this average. Dashed lines represent the 10% confidence intervals. Panel B of Figure 5 offers a similar plot for the 68 Type 2 buyers.[19]

Inspecting Panel A in Figure 5, one sees that, for sellers, the average proportion of first-order stochastically dominated actions starts around 35% in the first replication, and declines to 25% in the last replication. This proportion is always significantly lower than 50%, so we reject the pure-noise hypothesis that participants select randomly and uniformly any quantity between 0 and 4. But the proportion of FOS dominated actions is significantly larger than zero, so we reject the rational-choice hypothesis. Similarly, Panel B in Figure 5 shows that, for buyers, the frequency of dominated actions starts from slightly less than 35% in the first replication and then steadily declines to less than 25% in the last replication. Again, both purely noisy and perfectly rational choices are rejected. Thus participants are not perfectly rational but learn to make more rational choices as they become more

[17]In this latter case, we test the null hypothesis that the quantity does not change with the change of endowments in the aggregate-risk treatment; i.e., we test that the difference between the sum of odd-period fixed effects and the sum of even-period fixed effects has mean zero. For buyers, the null hypothesis of no treatment effect is rejected with a p-value smaller than 1%.

[18] At price 60 or for Type 3 participants, there are no first-order stochastically dominated actions.

[19]First-order stochastically dominated actions can be detected without making parametric assumptions on preferences only for Type 2, not for Type 3.

experienced. The results for the robustness experiment are offered in Figure OA.2 in the Online Appendix and are consistent with the baseline experiment.

We also display in Panel C of Figure 5 the evolution of the proportion of small and large mistakes. We define a small (respectively, large) mistake as an action in the first-order stochastically dominated quadrants illustrated in Panels A–B of Figure 1 that differs from 2 by less (respectively, more) than one unit.[20] Panel C of Figure 5 shows that both large and small mistakes tend to decrease with experience, consistent with learning. It also shows that large mistakes are less likely to occur than small ones.

Further statistical evidence is offered in Table 3. This table displays estimates for probit regressions in which the dependent variable is the indicator that an action is first-order stochastically dominated.[21] We control for participants' ID fixed effects. The main explanatory variable of interest is a proxy for a participant's experience, equal to the number of replications in which she has already participated.

The regression constant in column (1) of Table 3 is significantly negative. Together with the fact that the coefficient of the number of replications is also negative, this confirms that participants are more likely to select nondominated than dominated actions. Indeed, the constant being negative indicates that, given the probit regression model, the probability that the action is first-order stochastically dominated is lower than 50% already in the first replication of the experiment. However, this constant is also not very negative, indicating that many deviations from rational choice are observed.

Table 3 also confirms that the proportion of first-order stochastically dominated actions decreases with experience. In all our specifications, the coefficient on the experience variable is indeed significantly negative, which suggests that subjects learn to select nondominated actions over time.

Another explanatory variable of interest is the indicator that the replication involved random pricing, added in column (3) of Table 3. Its estimated coefficient is not significantly different from zero, consistent with the hypothesis that participants behave in the same way in the call and the random mechanisms, as already discussed above.

To further analyze learning in our experiment, the fourth specification in column (4) of Table 3 involves three additional regressors. First, the indicator that the price was previously

[20] As mentioned above, in Treatment II we cannot identify buyers' first-order stochastically dominated actions. So in Figure 5 we consider replications of Treatment I only.

[21] Our data include a total of 12,287 observations: one per participant, per replication, and per price, except that for one replication the data was lost because of a technical problem. We focus on Type 1 and Type 2 participants and on prices different from 60, because these are the cases in which we can identify first-order stochastically dominated actions. This leaves us with 8,110 observations (this corresponds to $68 \times 8 \times 10 + 68 \times 4 \times 10 = 8,160$ observations, minus the data we lost, five Type 1 participants placing orders at ten prices, which yields $8,160 - 50 = 8,110$ observations). We further drop 520 observations corresponding to the eight participants who never selected a dominated action. We thus end up with 7,590 observations.

selected as a market-clearing price is not significant, suggesting that participants do not pay more attention to prices at which they previously transacted. The two next variables are designed to capture different types of adaptive learning (Camerer and Ho (1999)). The indicator that the participant already selected a dominated action at that price displays a coefficient that is significantly positive, suggesting inertia in mistakes.[22] The indicator that the participant previously selected a dominated action at that price and incurred a loss as a result is not significant, suggesting past underperformance does not deter from selecting dominated actions.[23]

Despite the introduction of these regressors, the coefficient of the number of replications to which a subject has already participated remains significantly negative. This indicates that familiarity with the experimental setting helps subjects avoid dominated actions.

4.4 Market Size and Efficiency

Aggregate outcomes obtained when pooling all participants are close to competitive equilibrium and rational choice, but individual actions deviate from rationality. To analyze how market aggregation transforms individual mistakes into equilibrium outcomes, we study how the number of participants in a market influences the efficiency of prices and allocations. We focus on Treatment I with no aggregate risk, for which Implication 4 states the equilibrium price is equal to the expected dividend of 60.

We simulate markets by bootstrapping. For a market of size $2N$, we resample the original data set with replacement by randomly selecting the supply or demand schedules in one replication for N Type 1 sellers and N Type 2 buyers. For each N, we create $1,000$ simulated markets in which we aggregate demand and supply schedules to compute the market-clearing price. To document the impact of the number of market participants on pricing efficiency, Figure 6 plots the distribution of market-clearing prices when N varies from 5 to 80, and correspondingly the total number of participants in a market varies from 10 to 160.

When the market size is small, the distribution of market-clearing prices across the eleven prices in the grid is quite flat, even though the mode is at $S = 60$. This indicates that noise in individual demand and supply schedules translates into frequent pricing inefficiencies. When the market size increases, the distribution of the market clearing price becomes more and more concentrated around $S = 60$. To illustrate, consider what happens when the number of participants goes from 10 to 160. For 10 participants, the frequency with which the market clears at $S = 60$ is less than 15%. For 160 participants, that frequency rises to 50%.

[22]This is in contrast with the Law of Simulated Effect, according to which participants stop playing suboptimal actions when comparing them to counterfactuals.

[23]This is in contrast with the Law of Actual Effect, according to which agents adjust their behavior based on past realized outcomes.

To study the impact of the number of participants on allocative efficiency, we also use the bootstrap simulated markets, and we rely on the implication from theory that, without aggregate risk, risk averse participants should be fully hedged in equilibrium, i.e., they should trade 2 shares. To measure allocative efficiency, we compute, for each participant in each simulated market, the absolute difference between the quantity traded and 2. We then average this absolute difference across market participants. Figure 7 plots the distribution of this measure of allocative efficiency in our simulated markets when the number of market participants varies from 10 to 160. In equilibrium, this measure should be zero. If it is larger than zero, this indicates that risk sharing is imperfect and that risk-averse market participants are leaving gains from trade on the table.

Figure 7 shows that, for all the quartiles considered, the amount of unexploited gains from trade decreases with the number of market participants. For 10 market participants, the median of our measure of allocative efficiency in the simulated markets is above 0.25, more than 12.5% of the optimal trade, i.e., 2. By contrast, for 160 participants, that median is close to 0. Therefore, participants' trades deviate from 2, but less so when the number of market participants is large.

In the next section, we propose and estimate a random-choice model of individual behavior that is consistent with these observations on price and allocative efficiency.

5 A Random-Choice Model: Theory and Tests

The above reported findings are puzzling. On the one hand, aggregate outcomes are in line with the implications of rational choice in competitive markets. On the other hand, a large fraction of individual actions are inconsistent with rational choice. The latter suggests the prevalence of bounded rationality, the former that individual boundedly rational choices somehow add up to "rational" aggregate outcomes. To reconcile these findings, we propose a bounded-rationality model of individual choice along the lines of Luce (1959)'s random-choice model, derive its theoretical predictions, and confront them to the data. We finally use this framework to quantify allocative efficiency in our experimental financial markets.

5.1 The Model

As in Luce (1959), we assume that, instead of deterministically selecting an optimal action, each participant randomly chooses among available actions, putting larger probability on actions yielding higher utility.[24]

[24]In our competitive setting, agents only need to condition on the price, and need not form expectations about the actions of the others. Hence, our framework is simpler than the quantal response equilibrium, in which "[p]layers choose among strategies [...] based on their relative expected utility [...] and assume other players do so as well." (McKelvey and Palfrey (1995, p. 6)).

For every participant i of type T, let $U_i^T(q,S)$ be her utility of trading quantity q at price S. The function is defined for all admissible quantities, $q \in [0,4]$, and prices, $S \in \mathcal{S} \equiv \{52,\dots,62\}$. We assume that $U_i^T(q,S)$ is continuous in q for all S.

We want to specify the distribution of actions so that the probability that participant i selects quantity q at price S is increasing in $U_i^T(q,S)$. To do so, we consider a continuous weight function $\Phi_i^T : \mathbb{R}^{\mathcal{S}} \to \mathbb{R}$ such that the density of a quantity schedule $(q(S))_{S\in\mathcal{S}} \in [0,4]^{\mathcal{S}}$ is given by

$$f_{U_i^T,\Phi_i^T}((q(S))_{S\in\mathcal{S}}) \equiv \frac{\Phi_i^T((U_i^T(q(S),S))_{S\in\mathcal{S}})}{\int_0^4 \dots \int_0^4 \Phi_i^T((U_i^T(q_S,S))_{S\in\mathcal{S}})\, \mathrm{d}q_{52}\dots \mathrm{d}q_{62}}. \tag{4}$$

This formulation is very general and allows for correlation between a participant's quantity choices at different prices. For each $S \in \mathcal{S}$, we denote by $f_{U_i^T,\Phi_i^T,S}$ the corresponding marginal pdf. For any quantities $q, q' \in [0,4]$, $f_{U_i^T,\Phi_i^T,S}(q) > f_{U_i^T,\Phi_i^T,S}(q')$ if and only if $U_i^T(q,S) > U_i^T(q',S)$. We thus have the property that, at any given price, quantities delivering a higher utility are more likely to be selected.

Finally, we assume that, from the experimenter's viewpoint, for every type T, the utility functions U_i^T and the weight functions Φ_i^T of the N participants of type T are realizations of iid random elements $(\tilde{U}_i^T,\tilde{\Phi}_i^T)_{i=1}^N$ with common distribution $\mathbf{P}^T$, and that, given their realized characteristics (U_i^T,Φ_i^T), these participants draw their quantity schedules according to f_{U_i,Φ_i} independently. Thus the quantity schedules of participants of type T are iid random vectors $\tilde{q}_i^T$ with a common distribution $\mathbf{Q}^T = \mathbf{P}^T \otimes f_{U^T,\Phi^T}$.

5.2 Theoretical Predictions and Tests

This subsection studies financial market outcomes when agents make their decisions using the random-choice model. It helps explaining how the prevalence of mistakes that we observe in the data at the individual level is consistent with the experimental validity of predictions from competitive equilibrium theory at the aggregate level. It then shows that the random-choice model generates distributions of mistakes that are in line with our experimental findings. All the proofs are provided in Appendix A.3.

5.2.1 A Convergence Theorem

Focusing on Treatment I, with no aggregate risk, we show that suitably normalized aggregate experimental market outcomes, both in terms of prices and quantities, converge to the competitive equilibrium outcome when the number of participants goes to infinity.

Theorem 1 *In Treatment I, if all participants are FOSD risk-averse and behave according to the random-choice model, then, with probability* 1, *when the number of participants goes* *to ∞, the market-clearing price of the stock converges to* 60 *and the average per-participant trade converges to* 2.

Theorem 1 rationalizes the empirical findings for Treatment I presented in Section 4. When participants are FOSD risk-averse agents who behave according to the random-choice model, individual behavior is noisy and involves dominated actions. Yet, as the number of participants grows large, the aggregate outcome converges to the rational-choice competitive-equilibrium outcome. In particular, the price minimizing the gap between supply and demand eventually settles down to the expected value of the stock's dividend and the average quantity traded by each participant at that price converges to the full-hedge trade.

The proof of Theorem 1 relies on two observations, relying on the discussion in Section 3.1, in particular formula (3) and Implications 1 and 2:

1. First, at $S = 60$, trading $2 - \chi$ and $2 + \chi$ yield the same utility for Type 1 and Type 2 agents, and that utility is decreasing in χ for risk-averse agents. Correspondingly, the distribution of the quantities supplied or demanded by a risk-averse Type 1 or Type 2 participant at $S = 60$ is symmetric around 2, with a mode at 2. The strong law of large numbers then implies that the empirical average per-participant trade at price $S = 60$ converges to 2 with probability 1 as the number of participants goes to infinity. The gap between the average quantity supplied or demanded thus converges to zero.

2. Second, at $S \neq 60$, for Types 1 or Type 2 FOSD agents, the utility from trading $2 - \chi$ is strictly larger or strictly lower than the utility from trading $2 + \chi$. Consequently, the distribution of the quantities supplied or demanded is asymmetric. Moreover, the distributions for Type 1 sellers and Type 2 buyers are skewed in opposite ways. The strong law of large numbers then implies that, with probability 1, when the number of participants becomes large enough, the gap between the average quantity supplied or demanded is larger than zero.

Together, these two observations imply that, in Treatment I, as the number of participants goes to infinity, the price minimizing the gap between supply and demand eventually settles down to $S = 60$, and the average per-participant trade converges to 2.

5.2.2 Asymptotic Properties of the Distribution of Mistakes

We now turn to the distribution of first-order stochastically dominated actions, or mistakes. According to Implications 1 and 2, such mistakes can be identified in Treatment I for FOSD Type 1 and 2 participants at prices $S \neq 60$. We establish two large-sample implications of our random-choice model that shed light on our previous empirical findings. We first show that dominated actions in any quantity interval on one side of $q = 2$ tend to be less frequent than nondominated actions in the symmetric interval on the other side of $q = 2$. We next show that smaller mistakes tend to be more frequent than larger ones. For conciseness, we focus on the supply behavior of Type 1 participants at prices $S < 60$; the other cases lead to similar predictions.

Implication 1 states that, at prices $S < 60$, supplying less than 2 shares first-order stochastically dominates supplying more than 2 shares for Type 1 agents. Therefore, the random-choice model implies that, at price $S < 60$, Type 1 agents will more often choose to trade less than 2 shares than more than 2 shares, i.e.,

$$f_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2-\chi) > f_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2+\chi), \quad \chi \in (0,2].$$

This implies

$$F_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2+\bar{x}) - F_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2+\underline{x}) < F_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2-\underline{x}) - F_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(2-\bar{x}), \quad 0 \le \underline{x} < \bar{x} \le 2,$$

where $F_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}$ is the cdf corresponding to $f_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}$. That is, the probability, under the random-choice model, that i selects a quantity in the nondominated interval $[2-\overline{x}, 2-\underline{x}]$ is higher than the probability that she selects a quantity in the dominated interval $[2+\underline{x}, 2+\overline{x}]$, symmetric with respect to 2. The following convergence result is then a consequence of the Glivenko–Cantelli theorem.

Theorem 2 *In Treatment I, if all Type 1 participants are FOSD and behave according to the random-choice model, then, for all* $S < 60$ *and* $\Delta x \equiv \overline{x} - \underline{x} \in (0,2)$, *when the number of participants goes to* ∞, *with probability* 1 *the fraction of Type 1 trades in any nondominated interval* $[2-\overline{x}, 2-\underline{x}] \subset [0,2)$ *is larger than in the dominated interval* $[2+\underline{x}, 2+\overline{x}] \subset (2,4]$.

Similar implications hold at any price $S > 60$ and for Type 2 participants. Theorem 2 implies that, in a large population of FOSD Type 1 agents who behave according to the random-choice model, the distribution of supply tends to be asymmetric around $q = 2$ at any price $S \neq 60$, and, more generally, in any two symmetric intervals of a given length on both sides of $q = 2$. This implies that nondominated actions, on one side of $q = 2$, should be more frequent than dominated actions, on the other side of $q = 2$. This implication is in line with what we observe in our experiment, as illustrated in Figure 2.

Table 4 offers a formal test of the implications of Theorem 2. Panel A shows that we can reject the hypothesis that actions are equally likely to be dominated or not. Instead, the test favors the hypothesis that dominant actions are more likely than dominated ones. This shows that the distribution of actions is asymmetric. These results obtain when using all the data, as well as when focusing on quantities close to 2 or far from 2.

At price $S = 60$, the random-choice model implies (irrespective of whether agents are risk-averse or not) that the distribution of quantities is symmetric around its mean of 2, because, for each $\chi \in (0, 2]$, trading a quantity $2 + \chi$ leads to the same lottery over final wealth as trading a quantity $2 - \chi$. This leads to the following corollary.

Corollary 1 *In Treatment I, if all Type 1 participants behave according to the random-choice model, then, at* $S = 60$ *and for each* $\Delta x = \overline{x} - \underline{x} \in (0, 2)$, *with probability* 1, *when the number of participants goes to* ∞, *these Type 1 participants empirically select as often quantities in any interval* $[2 - \overline{x}, 2 - \underline{x}] \subset [0, 2)$ *as in the symmetric interval* $[2 + \underline{x}, 2 + \overline{x}] \subset (2, 4]$.

A similar implication holds for Type 2 participants. Table 4, Panel B shows that, at price $S = 60$ we cannot reject the hypothesis that quantity offers in the interval $(0, 1)$ are as likely as quantity offers in the symmetric interval $(3, 4)$, and similarly for the intervals $[1, 2)$ and $(2, 4]$. This is consistent with the symmetry property of the empirical distribution of actions established in Corollary 1

Finally, we explore the distribution of mistakes under the random-choice model. Again, we focus on the supply behavior of Type 1 participants at prices $S < 60$. Consider two possible quantity choices $4 \geq \overline{q} > \underline{q} > 2$. At any price $S < 60$, selling $\underline{q}$ or $\overline{q}$ shares is a mistake from the perspective of an FOSD Type 1 agent. However, selling $\overline{q}$ is a worse mistake. Indeed, as shown in the proof of Implication 1, selling $\overline{q}$ shares leads to a final wealth equal to that obtained by selling $\underline{q}$ shares, minus a positive constant $(60 - S)(\overline{q} - \underline{q})$, plus a zero-mean lottery

$$\left(-60(\overline{q} - \underline{q}), \frac{1}{2}; 60(\overline{q} - \underline{q}), \frac{1}{2}\right).$$

Hence at price $S < 60$, FOSD risk-averse Type 1 agents get more utility from trading $\underline{q}$ than from trading $\overline{q}$. The random-choice model then implies they should select $\underline{q}$ more often than $\overline{q}$. Our final result then follows along the same lines as Theorem 2.

Theorem 3 *In Treatment I, if all Type 1 participants are FOSD and behave according to the random-choice model, then, for all* $S < 60$ *and* $\Delta x \equiv \overline{x} - \underline{x} \in (0, 2)$, *with probability* 1, *when the number of participants goes to* ∞, *the fraction of Type 1 trades in any "small mistake" interval* $[2 + \underline{x}, 2 + \overline{x}]$ *is larger than in the "large mistake" interval* $[2 + \underline{x} + \eta, 2 + \overline{x} + \eta]$.

Again, similar implications hold at any price $S > 60$ and for Type 2 participants. This is consistent with the evidence presented in Panel C of Figure 5: small mistakes (supplying a quantity in $(2, 3]$ for a Type 1 participant at a price $S < 60$) are empirically more frequent than large mistakes (supplying a quantity in $(3, 4]$ at the same price). We offer a formal nonparametric test of the implications of Theorem 3 in Table 4, Panel A. Small mistakes (dominated actions near $q = 2$) appear significantly more frequent than large mistakes (dominated actions far from $q = 2$).

5.3 Quantifying Allocative Efficiency

In this section, we quantify the gains from trade in our experimental market. To do so, we rely on a structural estimation of the special case of our random-choice model with constant relative risk aversion and logit weighting.

5.3.1 Random Choice with Constant Relative Risk Aversion and Learning

Model Specification First, we assume that every participant i's preferences over state-contingent wealth have an expected-utility representation, for a Bernoulli utility function with constant relative risk aversion γ_i,

$$u_i(W) = \begin{cases} \frac{W^{1-\gamma_i}-1}{1-\gamma_i} & \text{if } \gamma_i \neq 1 \\ \ln(W) & \text{if } \gamma_i = 1 \end{cases}.$$

The expected utility of participant i of type T for given price S and quantity q writes as

$$EU_i^T(q,S) \equiv \frac{1}{2}\left[u_i(W^T(u,q,S)) + u_i(W^T(d,q,S))\right].$$

Next, we assume that participant i chooses quantities according to a logistic random-choice model, as in the quantal-response model of McKelvey and Palfrey (1995, 1998), for a multiplicative logit weight function with payoff responsiveness $\lambda_{i,n}$ in the n^{th} replication of the experiment,

$$\Phi_{i,n}((EU_S)_{S\in\mathcal{S}}) \equiv \prod_{S\in\mathcal{S}} \exp(\lambda_{i,n} EU_S).$$

That is, the quantity choices of every participant i of type T are independent across prices and the marginal pdf of her quantity choices at price S is given by

$$f_{EU_i^T,\Phi_{i,n},S}(q(S)) = \frac{\exp(\lambda_{i,n} EU_i^T(q(S),S))}{\int_0^4 \exp(\lambda_{i,n} EU_i^T(q,S))\,\mathrm{d}q}.$$

Finally, to capture potential learning effects while allowing for heterogeneity across subjects' payoff responsiveness, we assume that every participant i's payoff responsiveness in the n^{th} replication of the experiment is of the form

$$\lambda_{i,n} = \lambda_i + \delta_i(n-1).$$

Parameter λ_i captures the participant i's payoff responsiveness in the first replication, whereas parameter δ_i captures the speed at which she learns to make more rational choices.

Estimation We estimate by maximum likelihood the parameters $(\gamma_i, \lambda_i, \delta_i)$ for each seller i of Type 1 and for each buyer i potentially switching from Type 2 to Type 3, using 88 observations (8 periods $\times$ 11 prices).[25] We also estimate the model under the constraint that

[25] To estimate the model, we discretize the choice of quantities in $[0,4]$, using intervals of 0.05 length, as well as the choice of $(\gamma_i, \lambda_i, \delta_i)$ in $[-0.2,3] \times [0,50] \times [-5,5]$, using intervals of 0.1 length.

$\delta_i = 0$ (no learning), and again under the constraints that $\lambda_i = \infty$ and $\delta_i = 0$ (perfectly rational behavior).[26]

We report the average and quartiles of the distribution of parameters across participants in Table 5. The average and median coefficients of relative risk aversion in the unconstrained model are 0.61 and 0.80, respectively, consistent with the estimates obtained in the literature (see, for instance, Holt and Laury (2002)). The median payoff responsiveness is 0.70, also in line with previous estimates (see, for instance, Rogers et al. (2009)). The average responsiveness is much higher, due to the fact that 25% of participants appear to be close to perfect rationality, with a payoff responsiveness higher than 50. The median learning parameter is 0.10, in line with previous estimates (Rogers et al. (2009)). The first and the third quartiles of the distribution are 0.0 and 5.0, respectively, indicating large differences in participants' ability to learn.

We use a likelihood-ratio test to compare the goodness of fit of the constrained models, without learning and with perfect rationality, to the one of the unconstrained model with imperfect rationality and learning. We reject perfect rationality (p-value < 0.01) as well as the absence of learning (p-value $= 0.016$). This is in line with the results obtained without parametric assumption on preferences and choices.

Overall, heterogeneity among participants is significant, as shown by the relatively large discrepancy between the first and the third quartiles of parameter estimates. This indicates major differences across participants' risk aversion, rationality, and speed of learning. We also find positive correlation across participants between initial payoff responsiveness λ_i and speed of learning δ_i, as their correlation coefficient is equal to 0.34 (p-value < 0.01). Hence participants who initially tend to play nondominated actions more often than others also tend to learn faster than others. This suggests that both dispositions may be driven by a common factor.

5.3.2 Allocative Efficiency

Relying on the above presented CRRA-logit model and estimates, we compute several counterfactuals and use them to quantify allocative efficiency in our experimental market.

1. Given the estimated relative risk aversion of participant i, we compute the optimal supply or demand schedule of this participant, which we denote by $(q_i^*(S))_{S\in\mathcal{S}}$. In general, this optimal schedule differs from the schedule actually posted by participant i in the experiment, which we denote by $(q_i(S))_{S\in\mathcal{S}}$.

2. For each S, we also compute $EU_i(q_i^*(S), S)$, the expected utility of participant i from

[26] In the estimation, we implement an infinite payoff responsiveness by setting $\lambda_i = 50$. As a robustness check, we also tried an upper bound of 1,000 for λ_i; this did not significantly affect the log-likelihood.

trading her optimal quantity $q_i^*(S)$ at price S, and $EU_i(q_i(S), S)$, the expected utility of this agent from trading at price S according to the schedule she actually posted.

3. Given participants' optimal demand schedules in market m (corresponding to one trading round for one cohort), we can compute the theoretical price that would have cleared that market had they submitted their optimal schedules, which we denote by S_m^*. In general, this theoretical price differs from the price that actually cleared market m in the experiment, which we denote by S_m.

Based on these constructs, we first compute the proportion of participants who were not worse off participating in the market. To conduct this analysis, we focus on participants with nonnegative estimated coefficients of risk aversion. In market m we compute $EU_i(q_i(S_m), S_m)$, the expected utility of participant i in that cohort at the price S_m that prevailed in that market. We then compare this expected utility to its autarky counterpart, $EU_i(0)$, which corresponds to no trade and is thus independent of the price of the stock. We classify participant i, in that cohort, as not being made worse off by participating in the market if

$$EU_i(q_i(S_m), S_m) \geq EU_i(0). \tag{5}$$

For every market m, we compute the proportion $\Pi_m(S_m)$ of participants i who satisfy (5). To disentangle what stems from the price and what stems from individual decisions at a given price, we also evaluate that proportion at the rational competitive price S_m^* that would clear the market given the optimal schedules $(q_i^*(S))_{S \in \mathcal{S}}$ of participants in that market. For each market m, we compute $\Pi_m(S_m)$ and $\Pi_m(S_m^*)$. Then we compute the average of these proportions across markets. Notice that, by construction, the function Π_m is unaffected by increasing transformations of participants' expected utility.

Our empirical results on the proportion of participants who were not made worse off by participating in the market are presented in Table 6. In Treatment I, by construction, all trades at the rational competitive price $S_m^* = 60$ make risk-averse agents weakly better off than in autarky. So in Treatment I, $\Pi_m(S_m^*) = 100\%$. Actual prices, however, differ from 60. Indeed, in Treatment I, the average price in the experiment was 59.06. For these actual prices, some participants were worse off participating in the market, because they traded too much. As a result, the proportion $\Pi_m(S_m)$ was lower than 100%. For Type 1 participants, it was 89%, whereas, for Type 2 participants, it was 92%.

In Treatment II, there is aggregate risk and the price factors in a risk premium. Hence the rational competitive price varies with participants' risk aversion. On average, in our experiments, the rational competitive price in Treatment II was 58.61, but the actual price set in the experiments was lower, on average equal to 57.26. The proportion of Type 1 participants not made worse off by participating in the market was not very different from that in Treatment I: evaluated at the rational competitive price benchmark, it was 97%, whereas, at the lower price actually set in the experiments, it was 86%. For Type 3 participants, by contrast, the proportion of participants not made worse off by participating in the market was much lower: evaluated at the benchmark rational competitive price it was 42%, and at the price actually set in the experiment it was 59%. This is because, for Type 3 participants, autarky is relatively attractive as they start with a riskless endowment. When these agents buy too much, and at prices that are too high, they end up with lower expected utility than in autarky.

Second, for participants who were not worse off participating in the market, we compute the fraction of potential gains from trade that they actually reaped. For participant i, at price S, the fraction of potential gains from trade actually reaped is

$$\Psi_i(S) \equiv \frac{EU_i(q_i(S), S) - EU_i(0)}{EU_i(q_i^*(S), S) - EU_i(0)}.$$

The numerator in this fraction is the difference between the expected utility from the participant's actual trade at price S and her autarky utility, whereas the denominator is the difference between the expected utility from the participant's optimal trade at price S and her autarky utility. Because we focus on participants who were not worse off participating in the market, the numerator is positive. Moreover, $\Psi_i(S) \in [0, 1]$ as $EU_i(q_i^*(S), S) \geq EU_i(q_i(S), S)$ by definition, and $\Psi_i(S) = 1$ when participant i fully realizes all potential gains from trade at price S. Notice that, by construction, the function Ψ_m is unaffected by increasing affine transformations of participants i's expected utility. As above, we compute the fraction for each participant in each market, both at the benchmark rational competitive price and at the price that actually cleared that market in the experiment. Then we compute the average of that fraction across participants and markets.

Our empirical results on the fraction of potential gains from trade actually reaped are also presented in Table 6. In Treatment I, Type 1 participants not made worse off by participating in the market would have realized 81% of the gains from trade if the price had been equal to the benchmark rational competitive price of 60, and they realized 80% of the gains from trade at the prices actually set in the experiment. The corresponding percentages for Type 2 participants are 82% and 85%, respectively.

In Treatment II, Type 1 participants not made worse off by participating in the market would have realized 80% of the gains from trade if the price had been equal to the benchmark rational competitive price, and they realized 76% of the gains from trade at the prices actually set in the experiment. For Type 3 participants, the fraction of the gains from trade realized was much lower. At the benchmark rational competitive price (which was on average 58.61) they would have realized 31% of the potential gains from trade. At the prices actually set in the experiment (which was on average 57.26), they realized 43% of the potential gains from trade. Again, these low gains from trade reflect suboptimal demand schedules that induced excessive trades.

6 Conclusion

This paper experimentally tests rational-choice competitive equilibrium in complete financial markets. Our experiment is designed to closely emulate standard competitive equilibrium. Individual supply and demand schedules are elicited. In the benchmark market structure, supply and demand curves are aggregated, and crossed to set market-clearing prices. Comparing this benchmark to a setting in which the price is randomly set, we test, and cannot reject, the hypothesis that participants behave competitively. Moreover, the experiment is designed so that first-order and second-order stochastic dominance yield precise testable implications on individual behavior and market outcomes that hold for a very large class of preferences.

We find that around 25% of individual participants' actions are first-order stochastically dominated. Yet, when pooling all data, aggregate experimental outcomes are consistent with the predictions of rational-choice competitive equilibrium. Aggregate supply and demand cross at the expected dividend level when there is no aggregate risk, and at a lower price when there is aggregate risk, indicating the presence of a risk premium.

We develop a random-choice model, in the spirit of Luce (1959) and McKelvey and Palfrey (1995, 1998), that reconciles the apparently contradictory findings obtained at the aggregate and individual levels. Our random-choice model implies that individual deviations from rationality should average out as the number of market participants grows large, so that aggregate outcomes should converge to equilibrium. Simulations, based on bootstrapped experimental data, confirm that market clearing prices converge to equilibrium as the number of participants increases.

The random-choice model also implies that dominated actions should be less frequent than nondominated ones, and large mistakes less frequent than small ones. Our experimental data is in line with these predictions. Structural estimation of the special case of our random-choice model arising for CRRA utility and logit weighting enables us to quantify allocative efficiency. We find that around 80% participants are better off participating in the market than staying in autarky, whereas the remaining 20% are worse off.

In the end, our study suggests that, at the aggregate level, the predictions from rational choice competitive equilibrium can be quite robust to individual deviations from rationality. Therefore, we conjecture that market imperfections and incompleteness, rather that investors' irrationality, may drive empirical rejections of rational-choice competitive equilibrium asset pricing theory in the field. Extensions of the analysis to more complex settings and more comprehensive datasets would be useful to challenge and qualify this conjecture.

References

Arrow, Kenneth (1964) "The Role of Securities in the Optimal Allocation of Risk-Bearing," *Review of Economic Studies*, 31(2), 91–96.

Asparouhova, Elena, Peter Bossaerts, and John Ledyard (2024) "Price Formation in Multiple, Simultaneous Continuous Double Auctions, with Implications for Asset Pricing," Mimeo, University of Utah.

Asparouhova, Elena, Peter Bossaerts, Nilanjan Roy, and William Zame (2016) ""Lucas"' in the Laboratory," *Journal of Finance*, 71(6), 2727–2780.

Barber, Brad M. and Terrance Odean (2000) "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *Journal of Finance*, 55(2), 226–232.

——— (2001) "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116(1), 261–292.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964) "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9(3), 226–232.

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sébastien Pouget' (2005) "Judgemental Overconfidence, Self- Monitoring, and Trading Performance in an Experimental Financial Market," *Review of Economic Studies*, 72(2), 287–312.

Borch, Karl (1962) "Equilibrium in a Reinsurance Market," *Econometrica*, 30(3), 424–444.

Bossaerts, Peter, Leslie Fine, and John Ledyard (2002) "Inducing Liquidity in Thin Financial Markets through Combined-Value Trading Mechanisms," *European Economic Review*, 46(9), 1671–1695.

Bossaerts, Peter, Debrah Meloso, and William Zame (2015) "Dynamically Complete Experimental Asset Markets," Mimeo, University of Utah.

Bossaerts, Peter and Charles Plott (2004) "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," *Review of Finance*, 8(2), 135–169.

Bossaerts, Peter, Charles Plott, and William Zame (2007) "Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments," *Econometrica*, 75(4), 993–1038.

Breeden, Douglas T. (1979) "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7(3), 265–296.

Camerer, Colin and Teck Hua Ho (1999) "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67(4), 827–874.

Campbell, John Y. (2003) "Consumption-Based Asset Pricing," *Handbook of the Economics of Finance*, Volume 1, 803–887.

Crockett, Sean, John Duffy, and Yehuda Izhakian (2019) "An Experimental Test of the Lucas Asset Pricing Model," *Review of Economic Studies*, 86(2), 627–667.

Crockett, Sean, Daniel Friedman, and Ryan Oprea (2021) "Naturally Occurring Preferences and General Equilibrium: A Laboratory Study," *International Economic Review*, 62(2), 831–859.

DeBondt, Werner F.M. and Richard Thaler (1985) "Does the Stock Market Overreact?" *The Journal of Finance*, 40(3), 793–805.

Debreu, Gérard (1959) *Theory of Value*: *An Axiomatic Analysis of Economic Equilibrium*: New Haven: Yale University Press.

Fattinger, Felix (2021) "Trading Complex Risks," Mimeo, Vienna University of Economics and Business.

Fischbacher, Urs (2007) "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 10(2), 171–178.

Geanakoplos, John D. and Heraklis M. Polemarchakis (1986) "Existence, Regularity and Constrained Suboptimality of Competitive Allocations when the Asset Market is Incomplete," *Essays in Honor of Kenneth J. Arrow*: *Uncertainty, Information, and Communication*, Volume 3, 65–95.

Goeree, Jacob K. and Charles A. Holt (2004) "A Model of Noisy Introspection," *Games and Economic Behavior*, 46(2), 365–382.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), 247–272.

——— (2003) "Risk Averse Behavior in Generalized Matching Pennies Games," *Games and Economic Behavior*, 45(1), 97–113.

Grossman, Sanford J. and Joseph E. Stiglitz (1980) "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), 393–408.

Holt, Charles A. and Susan K. Laury (2002) "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644–1655.

Julliard, Christian and Anisha Ghosh (2012) "Can Rare Events Explain the Equity Premium Puzzle?" *The Review of Financial Studies*, 25(10), 3037–3076.

Kroll, Yoram and Haim Levy (1992) "Further Tests of the Separation Theorem and the Capital Asset Pricing Model," *American Economic Review*, 82(3), 664–670.

Kroll, Yoram, Haim Levy, and Amnon Rapoport (1988) "Experimental Tests of the Separation Theorem and the Capital Asset Pricing Model," *American Economic Review*, 78(3), 500–519.

Lucas, Robert E. (1978) "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–1445.

Luce, Duncan R. (1959) *Individual Choice Behavior*: *A Theoretical Analysis*: New York: Wiley.

Ludvigson, Sidney C. (2013) "Advances in Consumption-Based Asset Pricing: Empirical Tests," *Handbook of the Economics of Finance*, Volume 2, 799–906.

Machina, Mark J. (1982) ""Expected Utility" Analysis without the Independence Axiom," *Econometrica*, 50(2), 277–323.

Magnani, Jacopo, Jean Paul Rabanal, Olga A. Rud, and Yabin Wang (2022) "Efficiency of Dynamic Portfolio Choices: An Experiment," *Review of Financial Studies*, 35(3), 1279–1309.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995) *Microeconomic Theory*: New York: Oxford University Press.

McCabe, Kevin A., Stephen J. Rassenti, and Vernon L. Smith (1992) "Designing Call Auction Institutions: Is Double Dutch the Best?" *Economic Journal*, 102(410), 9–23.

McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6–38.

——— (1998) "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9–41.

Plott, Charles R. and Kirill Pogorelskiy (2017) "Call Market Experiments: Efficiency and Price Discovery through Multiple Calls and Emergent Newton Adjustments," *American Economic Journal*: *Microeconomics*, 9(4), 1–41.

Quiggin, John (1982) "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3(4), 323–343.

Rietz, Thomas A. (1988) "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22(1), 117–131.

Rogers, Brian W., Thomas R. Palfrey, and Colin F. Camerer (2009) "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," *Journal of Economic Theory*, 144(4), 1440–1467.

Rubinstein, Mark (1976) "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics*, 7(2), 407–425.

Segal, Uzi and Avia Spivak (1990) "First Order versus Second Order Risk Aversion," *Journal of Economic Theory*, 51(1), 111–125.

Shiller, Robert J. (2000) *Irrational Exuberance*: Princeton University Press.

Smith, Vernon L., Arlington W. Williams, W. Kenneth Bratton, and Michael G. Vannoni (1982) "Competitive Market Institutions: Double Auctions vs. Sealed Bid-Offer Auctions," *American Economic Review*, 72(1), 58–77.

Appendix

A.1 The Underlying Complete-Market Environment

Consider a one-period complete-market environment with two states of nature $\omega = u, d$ and two assets, a stock paying a dividend $D(\omega)$ in state ω, with $D(u) > 1 > D(d)$ and a bond paying 1 in each state. An agent initially holds a portfolio $\theta_0 \equiv (\theta_{0,s}, \theta_{0,b})$ of stocks and bonds, and receives an additional income $I(\omega)$ in state ω. In the market, the agent can trade to her final holdings of stocks and bonds, $\theta_1 \equiv (\theta_{1,s}, \theta_{1,b})$. Given security prices (p_s, p_b), the agent's budget constraint is $p_s(\theta_{1,s} - \theta_{0,s}) + p_b(\theta_{1,b} - \theta_{0,b}) \leq 0$. The agent maximizes her utility from state-contingent final wealth

$$W(\omega) \equiv I(\omega) + D(\omega)\theta_{1,s} + \theta_{1,b}, \quad \omega = u, d. \tag{A.1}$$

Assuming that the agent's utility is increasing with respect to her final wealth in each state, her budget constraint must be binding:

$$\theta_{1,b} = \theta_{0,b} - \frac{p_s}{p_b}\left(\theta_{1,s} - \theta_{0,s}\right). \tag{A.2}$$

Substituting (A.2) into (A.1) yields

$$W(\omega) = I(\omega) + \theta_{0,b} + D(\omega)\theta_{0,s} + \left[D(\omega) - \frac{p_s}{p_b}\right](\theta_{1,s} - \theta_{0,s}), \quad \omega = u, d. \tag{A.3}$$

Thus the only factor that affects the agent's final wealth in state ω, beyond her initial endowment $I(\omega) + \theta_{0,b} + \theta_{0,s}D(\omega)$, is the product of her net trade in the stock, $\theta_{1,s} - \theta_{0,s}$, by the profit margin on this trade, equal to the difference between the dividend and the price of the stock relative to that of the bond, $D(\omega) - \frac{p_s}{p_b}$. Therefore, because there are only two states of nature, the simple market structure in our experimental setting, in which participants can only trade the stock for the bond, implements a complete market structure. Letting $S \equiv \frac{p_s}{p_b}$ and $q \equiv |\theta_{1,s} - \theta_{0,s}|$, applying (A.3) to Type 1, Type 2, and Type 3 participants yields (A.4)–(A.5), (A.7)–(A.8), and (A.9)–(A.10), respectively.

A.2 Testable Implications of the Rational-Choice Model

Proof of Implication 1. By (A.3), if a Type 1 agent sells q shares at price S, then her final wealth levels in states u and d are

$$W^1(u, q, S) = 600 + (S - 120)q, \tag{A.4}$$

$$W^1(d, q, S) = 360 + Sq. \tag{A.5}$$

By (1)–(2) and (A.4)–(A.5),

$$\mu^1(q, S) = 480 + (S - 60)q \ \text{ and } \ \sigma^1(q, S) = 60\,|q - 2|. \tag{A.6}$$

(i) It follows from (A.6) along with our preliminary observations that, for each $\chi \in (0, 2]$, $L^1(2-\chi, S)$ first-order stochastically dominates $L^1(2+\chi, S)$ if $S < 60$, whereas $L^1(2+\chi, S)$ first-order stochastically dominates $L^1(2-\chi, S)$ if $S > 60$. This yields the first statement in (i). As for the second statement in (i), observe from (A.4)–(A.5) that, if $S < 60$ and $0 \leq \underline{q} < \overline{q} \leq 2$, then

$$W^1(u, \underline{q}, S) - W^1(u, \overline{q}, S) = (60 - S)(\overline{q} - \underline{q}) + 60(\overline{q} - \underline{q})$$
$$W^1(d, \underline{q}, S) - W^1(d, \overline{q}, S) = (60 - S)(\overline{q} - \underline{q}) - 60(\overline{q} - \underline{q}).$$

Hence $L^1(\underline{q}, S)$ is equal to $L^1(\overline{q}, S)$, plus a positive constant $(60-S)(\overline{q}-\underline{q})$, plus a zero-mean lottery $(-60(\overline{q} - \underline{q}), \frac{1}{2}; 60(\overline{q} - \underline{q}), \frac{1}{2})$. Thus, at $S < 60$, an FOSD risk-loving or risk-neutral Type 1 seller prefers the lottery $L^1(0, S)$ to any lottery $L^1(q, S)$, $q \in (0, 4]$, and similarly, at $S > 60$, she prefers the lottery $L^1(4, S)$ to any lottery $L^1(q, S)$, $q \in [0, 4)$. This yields the second statement in (i).

(ii) By (A.6), at $S = 60$, a Type 1 agent is perfectly hedged if she sells 2 shares, as $\sigma^1(2, 60) = 0$. More generally, if $0 \leq |q-2| < |q'-2| \leq 2$, then $L^1(q', 60)$ is a mean-preserving spread of $L^1(q, 60)$. Hence it is uniquely optimal for a risk-averse Type 1 agent to supply exactly 2 shares at $S = 60$. By contrast, a risk-loving Type 1 agent finds it optimal to supply 0 or 4 shares. Finally, a risk-neutral Type 1 agent is indifferent between supplying any quantities in $[0, 4]$ at price $S = 60$. The result follows. ■

Proof of Implication 2. By (A.3), if a Type 2 agent sells q shares at price S, then her final wealth levels in states u and d are

$$W^2(u, q, S) = 310 + (120 - S)q, \tag{A.7}$$
$$W^2(d, q, S) = 550 - Sq. \tag{A.8}$$

The proof then proceeds along similar lines as for Implication 1. The result follows. ■

Proof of Implication 3. By (A.3), if a Type 3 agent sells q shares at price S, then her final wealth levels in states u and d are

$$W^3(u, q, S) = 310 + (120 - S)q = 310 + (60 - S)q + 60q, \tag{A.9}$$
$$W^3(d, q, S) = 310 - Sq = 310 + (60 - S)q - 60 \times q. \tag{A.10}$$

Hence $L^3(q, S)$ is equal to Type 3 agent's safe endowment of 310, plus a constant $(60 - S)q$, plus a zero-mean lottery $(-60q, \frac{1}{2}; 60q, \frac{1}{2})$. The proof then proceeds along similar lines as for Implications 1–2. The result follows. ■

Proof of Implication 4. That there exists an equilibrium in which $S = 60$ and the N Type 1 and Type 2 participants all trade 2 shares, whereas the Type 3 participant, if any, trades 0 share, follows from Implications 1–3 for risk-averse agents. Now, suppose, by way of contradiction, that there exists an equilibrium with price $S > 60$. By Implication 1, FOSD Type 1 participants collectively supply at least $2N$ shares at S, whereas FOSD Type 2 participants collectively demand at most $2N$ shares at S, and the risk-averse Type 3 participant, if any, demands 0 share at S. Hence, it must be that each Type 1 and Type 2 agent trades exactly 2 shares at S. Hence, for every Type 2 participant i with utility $V_i^2(W(u), W(d))$ over state-contingent wealth, it must be that, for each $\chi \in (0, 2]$,

$$\begin{aligned} V_i^2(550 - 2S, 550 - 2S) &> V_i^2(310 + (2 - \chi)(120 - S), 550 - (2 - \chi)S) \\ &= V_i^2(550 - 2S + \chi(S - 120), 550 - 2S + \chi S) \end{aligned} \tag{A.11}$$

by (A.7)–(A.8). Now, the lottery $(S - 120, \frac{1}{2}; S, \frac{1}{2})$ has positive mean as $S > 60$. Hence, because participant i is assumed to be second-order risk-averse, it follows from Segal and Spivak (1990, Proposition 1) that, for $\chi > 0$ small enough,

$$V_i^2(550 - 2S + \chi(S - 120), 550 - 2S + \chi S) > V_i^2(550 - 2S, 550 - 2S),$$

in contradiction to (A.11). Thus there exists no equilibrium with price $S > 60$. A symmetric argument rules out an equilibrium with price $S < 60$. The result follows. ■

Proof of Implication 5. An equilibrium exists when all participants are FOSD risk-averse agents. Now, by Implication 1, FOSD Type 1 agents supply at least 2 shares at any price $S \geq 60$, whereas, by Implication 3, risk-averse Type 3 participants demand 0 share at any such price. Thus any equilibrium features a price $S < 60$. The result follows. ■

A.3 Theoretical Predictions of the Random-Choice Model

Proof of Theorem 1. The proof consists of two steps.

Step 1 Consider first what happens at price $S = 60$, at which every Type 1 participant i supplies a random quantity $\tilde{q}_i^1(60)$ given her characteristics $(\tilde{U}_i^1, \tilde{\Phi}_i^1)$. By Implication 1, as any such i is risk-averse, it is uniquely optimal for her to supply 2 shares at this price. That is, $\tilde{U}_i^1(2, 60) > \tilde{U}_i^1(q, 60)$ for all $q \in [0, 4] \setminus \{2\}$. By (4), this implies that

$$f_{\tilde{U}_i^1, \tilde{\Phi}_i^1, 60}(2) > f_{\tilde{U}_i^1, \tilde{\Phi}_i^1, 60}(q), \quad q \in [0, 4] \setminus \{2\}.$$

That is, $f_{\tilde{U}_i^1, \tilde{\phi}_i^1, 60}$ reaches its unique maximum at 2, which is thus the mode of the distribution of $\tilde{q}_i^1(60)$. Now, consider two possible quantity choices, $\underline{q} \equiv 2 - \chi$ and $\overline{q} \equiv 2 + \chi$, where $0 < \chi \leq 2$. When $S = 60$, $\underline{q}$ and $\overline{q}$ give rise to the same lottery. As a result, $\tilde{U}_i^1(\underline{q}, 60) = \tilde{U}_i^1(\overline{q}, 60)$ and thus, as above,

$$f_{\tilde{U}_i^1, \tilde{\Phi}_i^1, 60}(2 - \chi) = f_{\tilde{U}_i^1, \tilde{\Phi}_i^1, 60}(2 + \chi), \quad 0 < \chi \leq 2.$$

That is, the pdf $f_{\tilde{U}_i^1,\tilde{\Phi}_i^1,60}$ is symmetric around its mode. We conclude that

$$\mathbf{E}_{f_{\tilde{U}_i^1,\tilde{\Phi}_i^1,60}}[\tilde{q}_i^1(60)] = \int_0^4 q f_{\tilde{U}_i^1,\tilde{\Phi}_i^1,60}(q)\,\mathrm{d}q = 2$$

and, therefore,

$$\mathbf{E}_{\mathbf{Q}^1}[\tilde{q}_i^1(60)] = \int \mathbf{E}_{f_{\tilde{U}_i^1,\tilde{\Phi}_i^1,60}}[\tilde{q}_i^1(60)]\,\mathbf{P}^1(\mathrm{d}(\tilde{U}_i^1,\tilde{\Phi}_i^1)) = 2$$

for all i. Because Type 1 participants' quantity schedules $(\tilde{q}_i^1)_{i\geq 1}$ are iid, the strong law of large numbers yields

$$\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\tilde{q}_i^1(60) = 2 \tag{A.12}$$

with probability 1. The same reasoning for Type 2 participants yields

$$\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\tilde{q}_i^2(60) = 2 \tag{A.13}$$

with probability 1. It follows from (A.12)–(A.13) that

$$\lim_{N\to\infty}\frac{1}{N}\left|\sum_{i=1}^{N}[\tilde{q}_i^2(60)-\tilde{q}_i^1(60)]\right| = 0 \tag{A.14}$$

with probability 1.

Step 2 Suppose, by way of contradiction, that, with positive probability, the empirical market-clearing price of the stock does not converge to 60 as N goes to ∞. Then, because the price grid $\mathcal{S}$ is finite, there exists a price $S\neq 60$ such that, with positive probability,

$$S\in\arg\min_{S'\in\mathcal{S}}\frac{1}{\tilde{N}_k}\left|\sum_{i=1}^{\tilde{N}_k}[\tilde{q}_i^2(S')-\tilde{q}_i^1(S')]\right| \tag{A.15}$$

for all $k\geq 1$ along a random subsequence $(\tilde{N}_k)_{k\geq 1}$. It follows from (A.14)–(A.15) that

$$\lim_{k\to\infty}\frac{1}{\tilde{N}_k}\left|\sum_{i=1}^{\tilde{N}_k}[\tilde{q}_i^2(S)-\tilde{q}_i^1(S)]\right| = 0 \tag{A.16}$$

with positive probability. Because Type 1 and Type 2 participants' quantity schedules $(\tilde{q}_i^1)_{i\geq 1}$ and $(\tilde{q}_i^2)_{i\geq 1}$ are iid, the strong law of large numbers yields, for each $T=1,2$,

$$\lim_{k\to\infty}\frac{1}{\tilde{N}_k}\sum_{i=1}^{\tilde{N}_k}\tilde{q}_i^T(S) = \mathbf{E}_{\mathbf{Q}^T}[\tilde{q}_1^T(S)] = \int_0^4 q\,\mathbf{E}_{\mathbf{P}^T}[f_{\tilde{U}_1^T,\tilde{\Phi}_1^T,S}(q)]\,\mathrm{d}q \tag{A.17}$$

with probability 1, where the second equality follows from Fubini's theorem. Thus, by (A.16)–(A.17), we have

$$\int_0^4 q\,\mathbf{E}_{\mathbf{P}^1}[f_{\tilde{U}_1^1,\tilde{\Phi}_1^1,S}(q)]\,\mathrm{d}q = \int_0^4 q\,\mathbf{E}_{\mathbf{P}^2}[f_{\tilde{U}_1^2,\tilde{\Phi}_1^2,S}(q)]\,\mathrm{d}q. \tag{A.18}$$

Suppose first that $S < 60$ and consider two quantity choices, $\underline{q} \equiv 2 - \chi$ and $\overline{q} \equiv 2 + \chi$, where $0 < \chi \leq 2$. As $S < 60$, it follows from the arguments leading to Implications 1–2 that selling $\underline{q}$ (respectively, buying $\overline{q}$) first-order stochastically dominates selling $\overline{q}$ (respectively, buying $\underline{q}$) for an FOSD Type 1 (respectively, an FOSD Type 2) participant. As a result, $\tilde{U}_1^1(\underline{q}, S) > \tilde{U}_1^1(\overline{q}, S)$ and $\tilde{U}_1^2(\underline{q}, S) < \tilde{U}_1^2(\overline{q}, S)$ and thus, as above,

$$f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 - \chi) > f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 + \chi), \quad 0 < \chi \leq 2, \tag{A.19}$$

$$f_{\tilde{U}_1^2, \tilde{\Phi}_1^2, S}(2 - \chi) < f_{\tilde{U}_1^2, \tilde{\Phi}_1^2, S}(2 + \chi), \quad 0 < \chi \leq 2. \tag{A.20}$$

for any two pairs of characteristics $(\tilde{U}_1^1, \tilde{\Phi}_1^1)$ and $(\tilde{U}_1^2, \tilde{\Phi}_1^2)$ of the FOSD participant 1 of Type 1 and Type 2, respectively. Integrating (A.19)–(A.20) over $(\tilde{U}_1^1, \tilde{\Phi}_1^1)$ and $(\tilde{U}_1^2, \tilde{\Phi}_1^2)$ with respect to $\mathbf{P}^1$ and $\mathbf{P}^2$, respectively, and computing average quantities under $\mathbf{E}_{\mathbf{P}^1}[f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(q)]$ and $\mathbf{E}_{\mathbf{P}^1}[f_{\tilde{U}_1^2, \tilde{\Phi}_1^2, S}(q)]$ then yields

$$\int_0^4 q\, \mathbf{E}_{\mathbf{P}^1}[f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(q)]\, \mathrm{d}q < 2 < \int_0^4 q\, \mathbf{E}_{\mathbf{P}^2}[f_{\tilde{U}_1^2, \tilde{\Phi}_1^2, S}(q)]\, \mathrm{d}q,$$

in contradiction to (A.18). A similar contradiction can be derived if $S > 60$. We conclude that, with probability 1, the empirical market-clearing price of the stock converges to 60 as N goes to ∞. As the price grid $\mathcal{S}$ is finite, this implies that the empirical market-clearing price of the stock is constant and equal to 60 as soon as $N \geq N_0((\tilde{U}_i^1, \tilde{\Phi}_i^1, \tilde{U}_i^2, \tilde{\Phi}_i^2)_{i \geq 1})$, where $N_0((\tilde{U}_i^1, \tilde{\Phi}_i^1, \tilde{U}_i^2, \tilde{\Phi}_i^2)_{i \geq 1}) < \infty$ depends on the sample paths of the participants' characteristics. It then follows from Step 1 that the empirical average per-participant trade converges to 2 with probability 1 as N goes to ∞. Hence the result. ■

Proof of Theorem 2. Define, for $0 \leq \underline{x} < \overline{x} \leq 2$,

$$\begin{aligned} &\Lambda_S^1(\underline{x}, \overline{x}) \\ &\quad = \mathbf{E}_{\mathbf{P}^1}\Big[F_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 - \underline{x}) - F_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 - \overline{x}) + F_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 + \underline{x}) - F_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 + \overline{x})\Big]. \end{aligned} \tag{A.21}$$

We have $\Lambda_S^1(\underline{x}, \underline{x}) = 0$ for all $\underline{x} \in [0, 2)$ and

$$\frac{\partial \Lambda_S^1}{\partial \overline{x}}(\underline{x}, \overline{x}) = \mathbf{E}_{\mathbf{P}^1}\Big[f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 - \overline{x}) - f_{\tilde{U}_1^1, \tilde{\Phi}_1^1, S}(2 + \overline{x})\Big] > 0$$

for all $\overline{x} \in (\underline{x}, 2]$. Hence $\Lambda_S^1(\underline{x}, \overline{x}) > 0$ for any such $\underline{x}$ and $\overline{x}$, and thus, by continuity and compactness,

$$\inf_{\underline{x} \in [0, 2 - \Delta x]} \Lambda_S^1(\underline{x}, \underline{x} + \Delta x) > 0 \tag{A.22}$$

for all $\Delta x \in (0, 2)$. Now, let $\hat{F}_{S,N}^1$ be the empirical distribution of the quantities selected by N Type 1 participants at price S, defined by

$$\hat{F}_{S,N}^1(x) \equiv \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{\tilde{q}_i^1(S) \leq x\}}, \quad x \in [0, 4],$$

and, for all $0 \leq \underline{x} < \overline{x} \leq 2$, let $\hat{\Lambda}^1_{S,N}(\underline{x}, \overline{x})$ be the empirical excess frequency of Type 1 participants who select quantities in $[2-\overline{x}, 2-\underline{x}]$ rather than in $[2+\underline{x}, 2+\overline{x}]$ at price S in a sample of N Type 1 participants,

$$\hat{\Lambda}^1_{S,N}(\underline{x}, \overline{x}) \equiv \hat{F}^1_{S,N}(2-\underline{x}) - \hat{F}^1_{S,N}(2-\overline{x}) + \hat{F}^1_{S,N}(2+\underline{x}) - \hat{F}^1_{S,N}(2+\overline{x}). \tag{A.23}$$

Now, the random variables $(\tilde{q}^1_i(S))_{i\geq 1}$ are iid, with common cdf $\mathbf{E}_{\mathbf{P}^1}\big[F_{\tilde{U}^1_1, \tilde{\Phi}^1_1, S}\big]$. Thus

$$\mathbf{P}\Big[\lim_{N\to\infty} \Big\|\hat{F}^1_{S,N} - \mathbf{E}_{\mathbf{P}^1}\Big[F_{\tilde{U}^1_1, \tilde{\Phi}^1_1, S}\Big]\Big\|_\infty = 0\Big] = 1 \tag{A.24}$$

by the Glivenko–Cantelli theorem, where $\| \cdot \|_\infty$ is the sup sorm. Because, by (A.21) and (A.23), we have

$$\Big|\hat{\Lambda}^1_{S,N}(\underline{x}, \overline{x}) - \Lambda^1_S(\underline{x}, \overline{x})\Big| \leq 4\Big\|\hat{F}^1_{S,N} - \mathbf{E}_{\mathbf{P}^1}\Big[F_{\tilde{U}^1_1, \tilde{\Phi}^1_1, S}\Big]\Big\|_\infty$$

for all $0 \leq \underline{x} < \overline{x} \leq 2$, we conclude from (A.24) that

$$\mathbf{P}\left[\lim_{N\to\infty} \sup_{\underline{x}\in[0,2-\Delta x]} \Big|\hat{\Lambda}^1_{S,N}(\underline{x}, \underline{x}+\Delta x) - \Lambda^1_S(\underline{x}, \underline{x}+\Delta x)\Big| = 0\right] = 1 \tag{A.25}$$

for all $\Delta x \in (0, 2)$. Finally, we have

$$\begin{aligned}\inf_{\underline{x}\in[0,2-\Delta x]} \hat{\Lambda}^1_{S,N}(\underline{x}, \underline{x}+\Delta x) \geq & \inf_{\underline{x}\in[0,2-\Delta x]} \Lambda^1_S(\underline{x}, \underline{x}+\Delta x) \\ & - \sup_{\underline{x}\in[0,2-\Delta x]} \Big|\hat{\Lambda}^1_{S,N}(\underline{x}, \underline{x}+\Delta x) - \Lambda^1_S(\underline{x}, \underline{x}+\Delta x)\Big|\end{aligned}$$

for all N, and thus

$$\begin{aligned}&\mathbf{P}\Big[\lim_{N\to\infty} \inf_{\underline{x}\in[0,2-\Delta x]} \hat{\Lambda}^1_{S,N}(\underline{x}, \underline{x}+\Delta x) > 0\Big] \\ &\quad \geq \mathbf{P}\left[\inf_{\underline{x}\in[0,2-\Delta x]} \Lambda^1_S(\underline{x}, \underline{x}+\Delta x) > \lim_{N\to\infty} \sup_{\underline{x}\in[0,2-\Delta x]} \Big|\hat{\Lambda}^1_{S,N}(\underline{x}, \underline{x}+\Delta x) - \Lambda^1_S(\underline{x}, \underline{x}+\Delta x)\Big|\right] \\ &\quad = 1,\end{aligned}$$

where the equality follows from (A.22) and (A.25). Hence the result. ■

Proof of Corollary 1. This follows along the same lines as the proof of Theorem 2, using the fact that $\Lambda^1_{60}(\underline{x}, \overline{x}) = 0$ for $0 \leq \underline{x} < \overline{x} \leq 2$ by symmetry of the pdf $f_{\tilde{U}^1_1, \tilde{\Phi}^1_1, 60}$ around 2. Hence the result. ■

Proof of Theorem 3. This follows along the same lines as the proof of Theorem 2. Hence the result. ■

Figure 1: **Theoretical Predictions on Individual Supply and Demand**

Figure 1 shows which quadrants are first-order stochastically dominated and examples of Type 1's supply, and Type 2's and Type 3' demands. The predictions for risk-averse and risk-loving or -neutral agents are also indicated.

Panel A: Type 1's supply

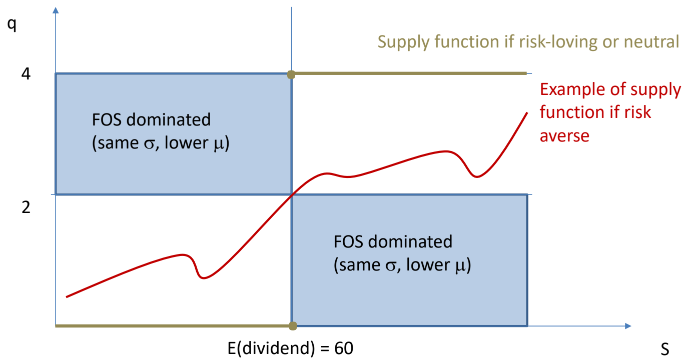

Panel B: Type 2's demand

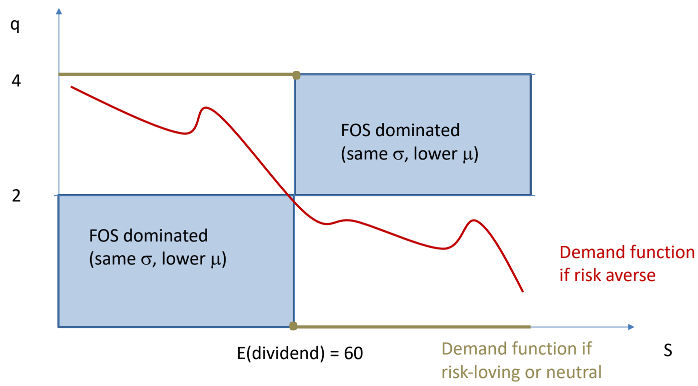

Panel C: Type 3's demand

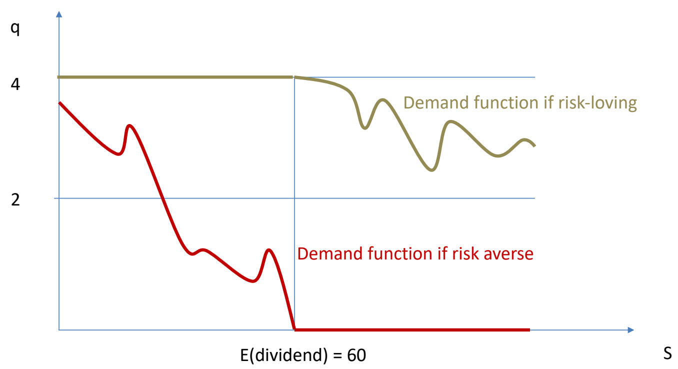

Figure 2: **Aggregate Supply and Demand**

Figure 2 illustrates the aggregate demand and supply schedules in our two treatments. There were 68 Type 1 sellers and 68 Type 2 buyers in Treatment I with no aggregate risk, and 68 Type 1 sellers and 73 Type 3 buyers in Treatment II with aggregate risk. We average the quantity supplied by sellers or demanded by buyers at each price across the four replications of the same treatment. We use the standard deviation across all observations of a subject's type in the same treatment to compute the 90% confidence intervals.

Panel A: Treatment I (no aggregate risk)

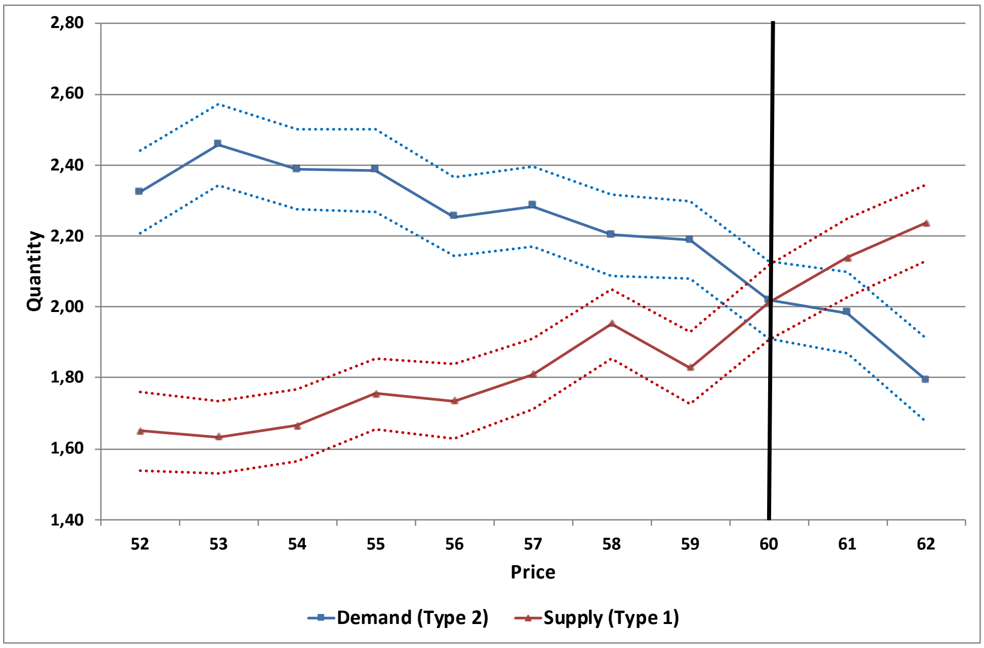

Panel B: Treatment II (aggregate risk)

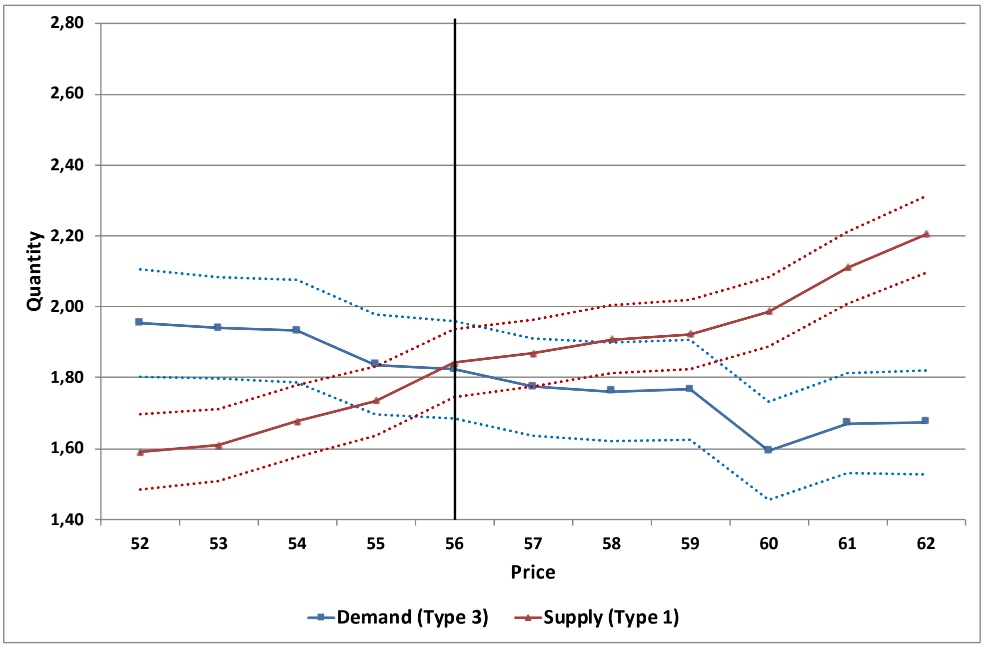

Figure 3: **Individual Supply and Demand Schedules**

Figure 3 illustrates the average demand and supply schedules submitted by participants depending on their type and on the treatment. There were 68 Type 1 sellers and 68 Type 2 buyers in Treatment I with no aggregate risk, and 68 Type 1 sellers and 73 Type 3 buyers in Treatment II with aggregate risk. We average the quantity supplied by sellers or demanded by buyers at each price across the four replications of the same treatment. Each line in grey corresponds to the demand or supply of one subject. Supply or demand schedules of a few subjects are highlighted in blue, red, or black.

Panel A: Type 1 sellers when there is no aggregate risk

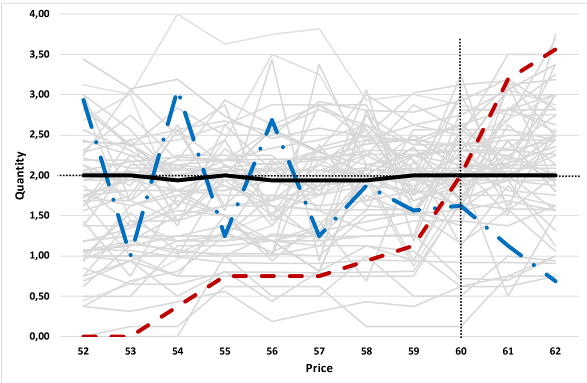

Panel B: Type 1 sellers when there is aggregate risk

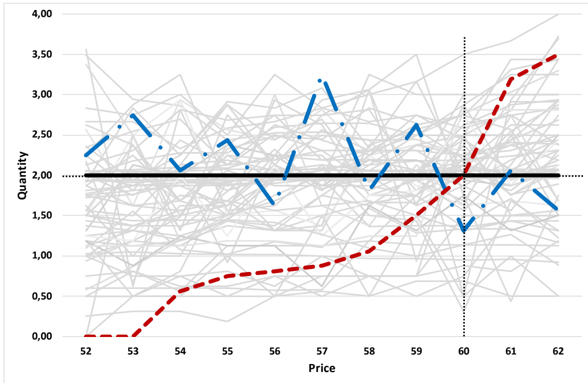

Panel C: Type 2 buyers when there is no aggregate risk

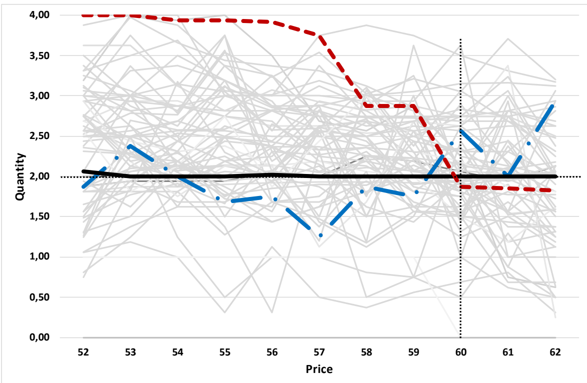

Panel D: Type 3 buyers when there is aggregate risk

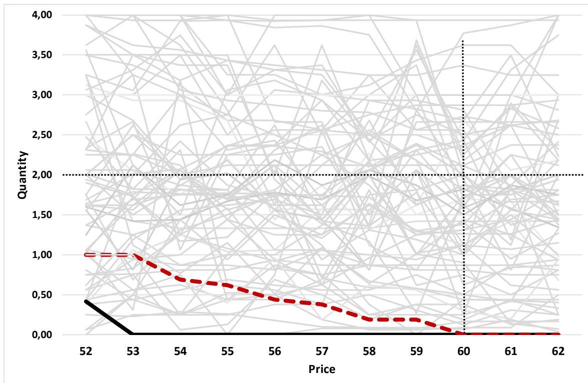

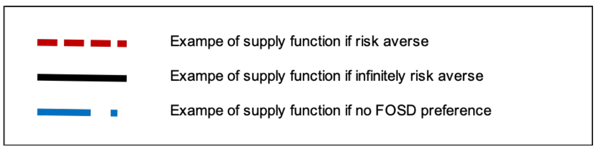

Figure 4: **Aggregate Supply and Demand in Call and Random Mechanisms**

Figure 4 illustrates the aggregate demand and supply schedules in our experiments in four cases, defined by the interaction between the treatment (no aggregate risk versus aggregate risk) and the determination of the price (call versus random mechanism). We average the quantity supplied by sellers or demanded by buyers at each price across the two replications of the same case.

Panel A: Treatment I (no aggregate risk)

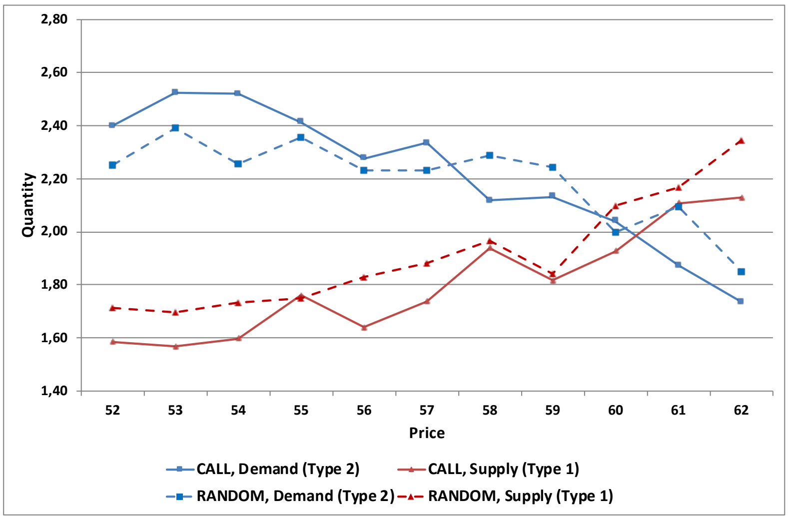

Panel B: Treatment II (aggregate risk)

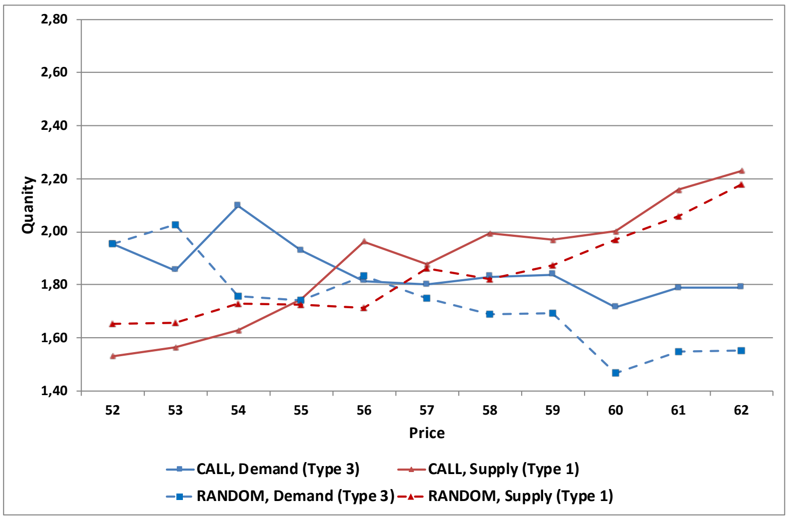

Figure 5: **Frequency of First-Order Stochastically Dominated Actions**

Figure 5 illustrates the evolution of deviations from first-order stochastic dominance by Type 1 sellers and Type 2 buyers. An action is first-order stochastically dominated for a Type 1 participant (respectively, a Type 2 participant) if she supplies a quantity $q > 2$ at a price $S < 60$ or a quantity $q < 2$ at a price $S > 60$ (respectively, if she demands a quantity $q < 2$ at a price $S < 60$ or a quantity $q > 2$ at a price $S > 60$). For each subject and each replication, we compute the proportion of first-order stochastically dominated actions for each relevant price. Panels A and B show, for each replication, the average of this proportion across the 68 subjects of Types 1 and 2, respectively. Panel C shows the evolution of the proportion of small and large mistakes, where a small (respectively, a large) mistake is defined as a quantity supplied or demanded in a dominated quadrant that differs from 2 by less (respectively, more) than one unit. We use the standard deviation across all observations of each type to compute the 90% confidence intervals.

Panel A: Type 1 sellers (eight replications)

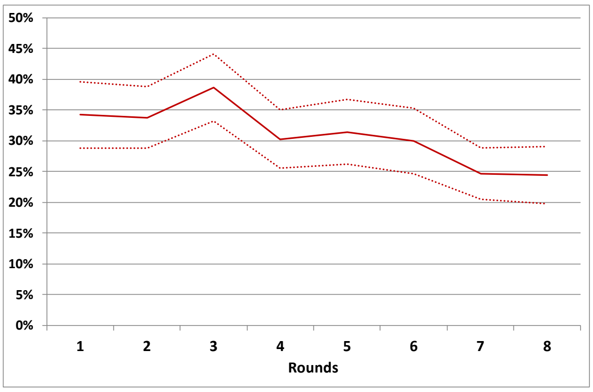

Panel B: Type 2 buyers (four replications)

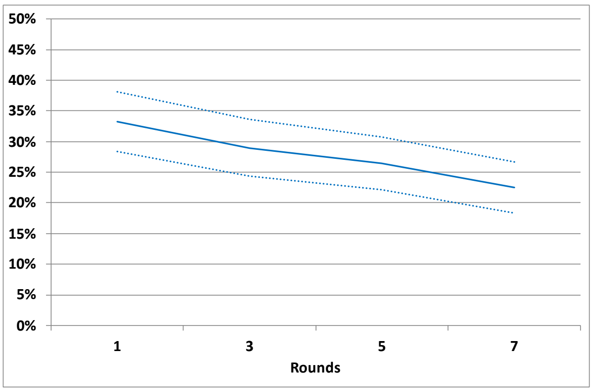

Panel C: Large and small mistakes for buyers and sellers

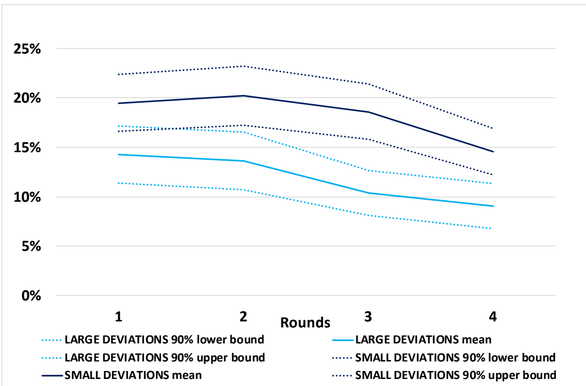

Figure 6: **Distribution of Market-Clearing Prices in Simulated Markets**

Figure 6 illustrates the impact of an increase in the number of participants on the distribution of market-clearing prices in Treatment I with no aggregate risk. We simulate markets by bootstrapping. For a market of size $2N$, we resample the original dataset with replacement by randomly selecting the supply or demand schedule in one replication for N Type 1 sellers and N Type 2 buyers. For each N and each of the $1,000$ simulated markets, we aggregate demand and supply schedules to compute the market-clearing price.

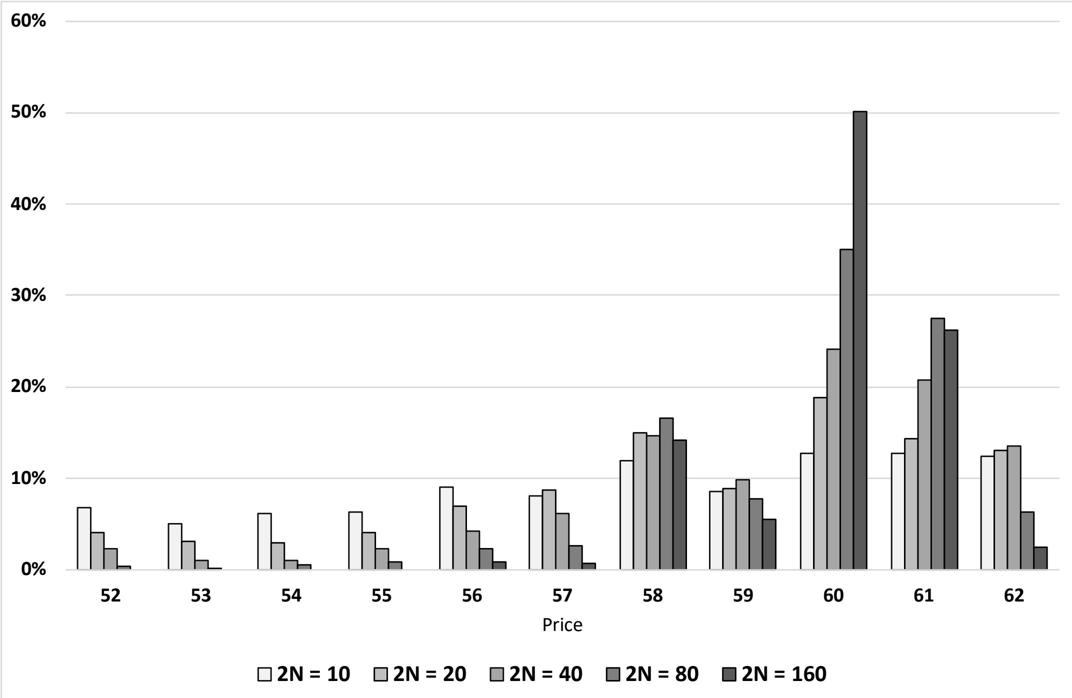

Figure 7: **Distribution of Allocative Efficiency in Simulated Markets**

Figure 7 illustrates the impact of an increase in the number of participants on allocative efficiency. Allocative efficiency is nonparametrically measured as the absolute difference between the quantity traded by a market participant and the quantity traded in equilibrium, averaged across market participants. We focus on Treatment I with no aggregate risk in which FOSD and risk aversion entails that market participants should be fully hedged in equilibrium and trade a quantity of 2. The measure of allocative efficiency is thus $\frac{1}{2N}\sum_i |q_i - 2|$, which should be zero in equilibrium. The figure shows the average of this measure for the different numbers of participants $2N$. We report the average within quartiles defined across the $1,000$ simulated markets. We simulate markets by bootstrapping. For a market of size $2N$, we resample the original dataset with replacement by randomly selecting the supply or demand schedule in one replication for N Type 1 sellers and for N Type 2 buyers. For each N and each of the $1,000$ simulated markets, we aggregate demand and supply schedules to compute the market-clearing price. Each market participant trades the quantity bid or offered at the market-clearing price.

Average distance to the theoretical quantity

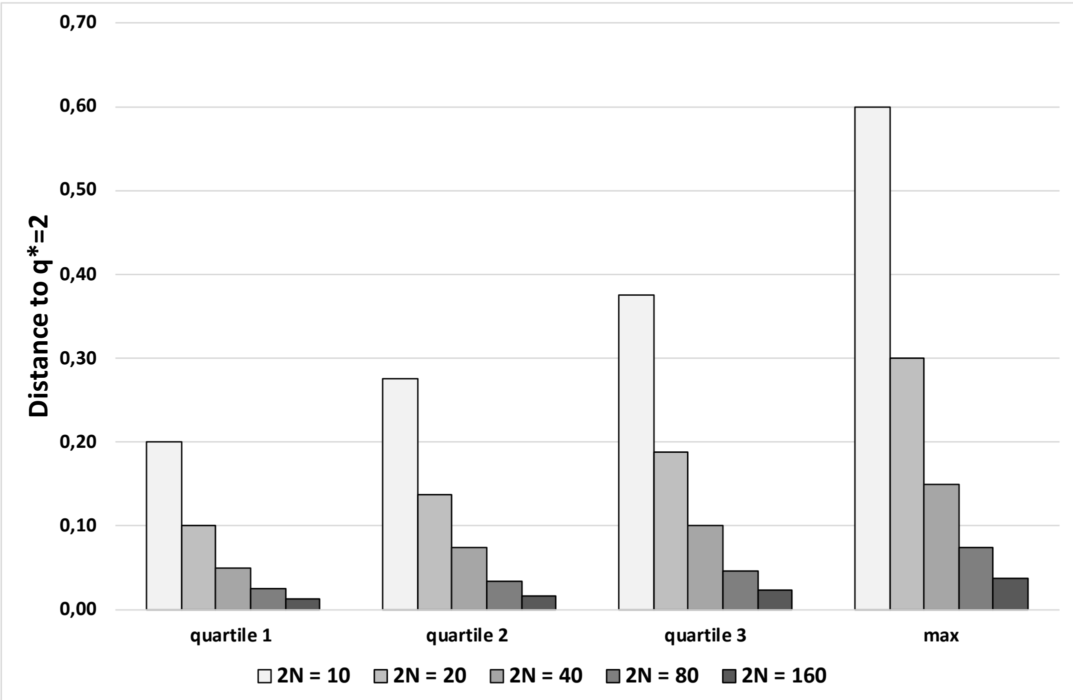

Table 1: **Experimental Protocol**

Our baseline experiment consists of eight sessions and includes in total 141 subjects who participate in eight replications. We alternate Treatment I in odd replications, consisting of Type 1 sellers and Type 2 buyers—plus possibly one Type 3 buyer when there is an odd number of subjects—and corresponding to no aggregate risk ("no AggR"), and Treatment II in even replications, consisting of Type 1 sellers and Type 3 buyers and corresponding to aggregate risk ("AggR"). For cohorts A, C, E and G, the first four replications involve random pricing, whereas in the last four replications the price is set to minimize the gap between supply and demand. For the other cohorts, in the first four replications prices are set to minimize the gap between supply and demand, whereas in the last four replications there is random pricing. Subjects are not informed of the treatment but they know which mechanism is used to set the price. Column "# part." indicates the number of participants in each session.

		Round							
Session	# part.	1	2	3	4	5	6	7	8
A	19	Random				Minimize gap Supply - Demand			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
B	11	Minimize gap Supply - Demand				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
C	19	Random				Minimize gap Supply - Demand			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
D	19	Minimize gap Supply - Demand				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
E	20	Random				Minimize gap Supply - Demand			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
F	20	Minimize gap Supply - Demand				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
G	15	Random				Minimize gap Supply - Demand			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
H	18	Minimize gap Supply - Demand				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR

Table 2: **Aggregate Risk, Random Mechanism, and Individual Demand and Supply Schedules**

Table 2 reports the estimates of a panel regression capturing the impact of the individual endowment and of random pricing on the quantity supplied and demanded, denoted by Q^{sell} and Q^{buy}, respectively. The explanatory variables are the price at which the quantity is demanded or supplied, an indicator that the replication involved aggregate risk, and an indicator that the replication involved random pricing. We exclude observations of Type 3 buyers who do not switch type in Treatment I with no aggregate risk. In columns (2) and (4), we include time fixed effects, thus excluding the redundant indicator that the replication involved aggregate risk. All regressions include participants' ID fixed effects. Standard errors are clustered by ID. The symbols ***, **, and * indicate significance at 1%, 5%, and 10% level, respectively, and t-statistics appear in parentheses.

	Quantity demanded or supplied			
	Q^{sell}	Q^{sell}	Q^{buy}	Q^{buy}
	(1)	(2)	(3)	(4)
Price	0.0368***	0.0582***	0.0187	−0.0451***
	(8.70)	(5.21)	(1.48)	(−4.54)
Aggregate Risk	0.0073		−0.4081***	
	(0.21)		(−3.26)	
Random Pricing	0.0405	0.0421	−0.0336	−0.0549
	(0.66)	(0.70)	(−0.54)	(−0.86)
# obs	5,929	5,929	5,929	5,929
R^2	0.8260	0.8292	0.7622	0.7719
Time FE	no	yes	no	yes
Test that the sum of Odd periods FE equals the sum of Even periods FE				
F-stat		0.10		11.83 ***
t-stat		0.32		−3.44***
p-value		0.7523		0.0010

Table 3: **Probit Regressions**

The dependent variable is an indicator that the action is first-order stochastically dominated. Regressions are estimated over 7,590 actions (see Footnote 21), by the 128 participants for whom we can identify dominated actions and who chose at least one dominated action. The explanatory variables include the number of replications in which the subject participated at the moment of the choice, an indicator that the replication involved random pricing, an indicator that trading already took place at that price, an indicator that the subject selected a dominated action at that price, and an indicator that this dominated action led to a loss-making trade. Standard errors are clustered by participant ID. Regressions (2)–(4) include participants' ID fixed effects. The symbols ***, **, and * indicate significance at 1%, 5%, and 10% level, respectively, and t-statistics appear in parentheses.

	First-order stochastically dominated action			
	(1)	(2)	(3)	(4)
Number of replications	-0.338^{***}	-0.0610^{***}	-0.0605^{***}	-0.0758^{***}
	(-2.82)	(-4.57)	(-4.60)	(-5.33)
Random pricing			0.0448	0.0473
			(0.93)	(0.97)
Trading already took place at that price				-0.0793
				(-1.21)
Subject selected dominated action at that price				0.6941^{***}
				(3.93)
Dominated action led to loss-making trade				-0.2750
				(-1.52)
Constant	-0.375^{***}			
	(-7.70)			
# obs	7,590	7,590	7,590	7,590
Pseudo R^2	0.0025	0.11	0.11	0.13
Participant fixed effects	no	yes	yes	yes

Table 4: **Distribution of Actions and Associated Tests**

This table displays the distribution of actions. Panel A refers to actions at prices different from 60. The proportion of actions that are first-order stochastically dominated and nondominated is calculated. For Type 1 participants, for example, nondominated actions correspond to selling a quantity $q < 2$ at prices $S < 60$: in this quadrant, the bound is 0, nondominated actions far from $q = 2$ are $0 < q \leq 1$ and those that are near $q = 2$ are $1 < q < 2$). The total number of actions is 1,360 in odd replications and 680 in even replication (except the 8th one for which it is 630 due to a technical issue that induced a data loss). Panel B corresponds to actions at a price of 60 at which there is no first-order stochastically dominated actions. The total number of actions is 136 in odd replications and 68 in even replication (except the 8th one for which it is 63 due to a technical issue that induced a data loss). In both panels, p-values are based on the binomial distribution with probability 0.5 under the null hypothesis. The number of trials correspond to the 8 replications, and the number of successes for H1 to the number of replications in which the data conforms with H1.

Panel A	Nondominated actions				Dominated actions		
(Price $S \neq 60$)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Replication	Bound	Far from $q = 2$	Near $q = 2$	$q = 2$	Near $q = 2$	Far from $q = 2$	Bound
1	171	252	282	196	203	172	84
2	49	135	153	113	119	86	25
3	157	220	269	254	214	136	110
4	51	127	157	139	108	63	35
5	155	173	324	314	200	130	64
6	53	100	169	154	111	66	27
7	215	176	326	322	175	76	70
8	75	83	157	161	87	45	22

Nonparametric tests Dominant vs dominated	Hypothesis H1	Successes H1	Trials	p-value under H0 (equally likely)
All	Prob[(1)+(2)+(3)] > Prob[(5)+(6)+(7)]	8	8	0.39%
Far from $q = 2$	Prob[(2)] > Prob[(6)]	8	8	0.39%
Near $q = 2$	Prob[(3)] > Prob[(5)]	8	8	0.39%
Bound	Prob[(1)] > Prob[(7)]	8	8	0.39%
Large vs Small	Prob[(5)] > Prob[(6)]	8	8	0.39%

Panel B	$q < 2$				$q > 2$		
(Price $S = 60$)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Replication	Bound	Far from $q = 2$	Near $q = 2$	$q = 2$	Near $q = 2$	Far from $q = 2$	Bound
1	19	18	22	37	19	18	3
2	3	9	8	14	17	10	7
3	11	14	15	29	27	23	17
4	5	8	9	15	20	9	2
5	9	7	25	37	25	24	9
6	5	7	9	19	21	4	3
7	14	12	18	38	31	14	9
8	6	8	12	19	10	7	1

Nonparametric tests $q < 2$ vs $q > 2$	Hypothesis H1	Successes H1	Trials	p-value under H0 (equally likely)
All	Prob[(1)+(2)+(3)] > Prob[(5)+(6)+(7)]	2	8	10.94%
Far from $q = 2$	Prob[(2)] > Prob[(6)]	2	8	10.94%
Near $q = 2$	Prob[(3)] > Prob[(5)]	2	8	10.94%
Bound	Prob[(1)] > Prob[(7)]	5	8	21.88%

Table 5: **Estimations of Random-Choice Models of Financial Markets**

We assume that every participant i has a Bernoulli utility function with constant relative risk aversion γ_i and chooses a quantity q at price S in the n^{th} replication with density $\frac{\exp(\lambda_{i,n} EU_i^T(q,S))}{\int_0^4 \exp(\lambda_{i,n} EU_i^T(x,S))\,\mathrm{d}x}$, where $EU_i^T(q,S)$ is the expected utility of her future wealth conditional on trading a quantity q and $\lambda_{i,n} = \lambda_i + \delta_i(n-1)$. Table 5 reports the descriptive statistics of the model's estimates on 141 participants. In Model (2), we impose $\delta_i = 0$. In Model (3), participants choose the quantity that maximizes their expected utility. Because the models are nested, the goodness of fit of a constrained model θ_c relative to the unconstrained model θ_0 is evaluated via the likelihood ratio $LR = -2[\ell(\hat{\theta}_0) - \ell(\hat{\theta}_c)]$, where ℓ is the log-likelihood of the model's estimation. Under the null hypothesis, the probability distribution of the log-likelihood-ratio statistic used to test nested models is approximated by a χ^2-squared distribution with degrees of freedom equal to the difference between the numbers of parameters in the two models.

(1) Unconstrained model

	Mean	Std Dev	Q1	Q2	Q3
γ	0.61	0.51	0.0	0.8	1
λ	16.98	22.29	0.1	0.7	50.0
δ	1.58	2.84	0.0	0.1	5.0
log$\mathcal{L}$	-371	46	-385	-374	-358

(2) Constrained model with $\delta = 0$ (no learning)

	Mean	Std Dev	Q1	Q2	Q3
γ	0.51	0.55	-0.1	0.5	0.9
λ	18.54	22.79	0.1	1.4	50.0
log$\mathcal{L}$	-374	44	-386	-376	-353
Goodness of fit relative to the unconstrained model					
LR				-5.79	
p-value				0.0161	

(3) Constrained model with $\lambda \to \infty$ (perfect rationality)

	Mean	Std Dev	Q1	Q2	Q3
γ	0.24	0.55	-0.1	-0.1	0.4
log$\mathcal{L}$	-1,248	199	-1,370	-1,337	-1,224
Goodness of fit relative to the unconstrained model					
LR				-1,753.81	
p-value				0.0000	

Table 6: **Allocative Efficiency**

This table offers various statistical measures of allocative efficiency. $\Pi_m(S)$ is the proportion of risk-averse participants in a cohort m whose trades translate in an increase in individual welfare compared to autarky, measured at a price S that is either the theoretical equilibrium price S_m^* or the realized price S_m. $\Psi_m(S)$ is the average fraction of the gains from trade extracted by these participants. We exclude observations of Type 3 buyers who do not switch type in Treatment I with no aggregate risk. The average and the standard deviation of the realized price are computed across the different cohorts of participants. The first column corresponds to statistics computed on all types, the second to Type 1 sellers, the third to Type 2 buyers (who are at risk before trading), and the fourth to Type 3 buyers (who are not at risk before trading).

	Replications			
No aggregate risk (Treatment I)	All types	Type 1 sellers	Type 2 buyers	Type 3 buyers
Theoretical price S_m^*	60.00	60.00	60.00	NA
$\Pi_m(S_m^*)$	100%	100%	100%	NA
$\Psi_m(S_m^*)$	81%	81%	82%	NA
Average realized price S_m	59.06	59.06	59.06	NA
Standard deviation of S_m	2.79	2.79	2.79	NA
$\Pi_m(S_m)$	89%	89%	92%	NA
$\Psi_m(S_m)$	83%	80%	85%	NA

	Replications			
Aggregate risk (Treatment II)	All types	Type 1 sellers	Type 2 buyers	Type 3 buyers
Theoretical price S_m^*	58.61	58.61	NA	58.61
$\Pi_m(S_m^*)$	77%	97%	NA	42%
$\Psi_m(S_m^*)$	69%	80%	NA	31%
Average realized price S_m	57.26	57.26	NA	57.26
Standard deviation of S_m	2.98	2.98	NA	2.98
$\Pi_m(S_m)$	78%	86%	NA	59%
$\Psi_m(S_m)$	68%	76%	NA	43%

Online Appendix

OA.1 Instructions for the Experiment

INSTRUCTIONS

You are going to participate in an economic experiment. Various research organizations have provided funds for conducting this research. If you follow the instructions below and make good decisions, you can earn a significant amount of money that will be paid to you in cash after the experiment.
In this experiment, you can trade a risky asset. The dividend offered by this asset is determined by a coin toss. Heads indicates that the dividend is 120 ECU (Experimental Currency Unit), and Tails indicates that the dividend is 0 ECU. The coin toss represents the risk that exists in the economic environment.

There are 11 potential prices: 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, et 70 ECU. If you hold assets, we will ask you what quantity you would like to sell at each potential price. If you do not hold assets, we will ask you what quantity you would like to buy at each potential price. You can choose any quantity between 0 and 4, included (within the technical limit of 8 decimal figures).

We will start by organizing 4 market sessions in which the transaction price will be determined by minimizing the distance between the aggregate supply (sum of all offers to sell) and demand (sum of all offers to buy). We will then organize 4 market sessions in which the transaction price will be randomly determined, all prices being equally likely.

For each of these 8 market sessions, you will start by being informed of your initial endowment of assets and of ECU and of your potential additional income (that also depends on the economic environment, i.e., of the result of the coin toss).

Then, at the end of each of the 8 market sessions, we will compute your final wealth that will depend on your initial endowments (in assets and in ECU), on your offer to buy or sell, on the transaction price, on the value of the dividend distributed by the risky asset and on the potential additional income determined by the coin toss.

The positions in assets and ECU are not transferred from session to session. The different market sessions are independent: you start each market session with a new situation and new endowments, and a new coin toss will be done.

Example 1: You start with 3 units of the risky asset, no ECU. You will also receive an additional income of 100 ECU if Heads comes up, and 300 ECU if Tails comes up. Suppose that the transaction price (whether determined randomly or by minimizing the distance between supply and demand) is 65 ECU and that, at this price, you offered to sell 2.5 units of the asset.

If Heads comes up, your final wealth is:

Proceeds from asset sale:	2.5 x65	= 162.5
+ Dividends on asset holdings:	(3 – 2.5) x120=	+ 60
+ Additional income if Heads:		+ 100
Total if Heads:		= 322.50 ECU

If Tails comes up, your final wealth is:

Proceeds from asset sale:	2.5 x65	= 162.5
+ Dividends on asset holdings:	(3 – 2.5) x0 =	+ 0

+ Additional income if Tails:		+ 300
Total if Tails:		= 462.50 ECU

Example 2: You start with no risky asset and 300 ECU. You will also receive an additional income of 120 ECU if Heads comes up, and 190 ECU if Tails comes up. Suppose that the transaction price (whether determined randomly or by minimizing the distance between supply and demand) is 55 ECU and that, at this price, you offered to buy 1.5 units of the asset.

If Heads comes up, your final wealth is:

Initial endowment in ECU:		300
- Cost of purchase:	1.5 x55 =	- 82.5
+ Dividends on asset holdings:	1.5 x120 =	+ 180
+ Additional income if Heads:		+ 120
Total if Heads:		= 517.5 ECU

If Tails comes up, your final wealth is:

Initial endowment in ECU:		300
- Cost of purchase:	1.5 x55 =	- 82.5
+ Dividends on asset holdings:	1.5 x0 =	+ 0
+ Additional income if Tails:		+ 190
Total if Tails:		= 407.5 ECU

In order to pay you, we will randomly select one of the first 4 market sessions and one of the last 4 market sessions. We will pay you the sum of the final wealth you obtained during the two selected sessions at the exchange rate of 10 ECU = 1 €. For example, suppose that your final wealth is 210 ECU at market session 1, 430 ECU at session 2, 70 ECU at session 3, 140 ECU at session 4, 380 ECU at session 5, 540 ECU at session 6, 80 ECU at session 7, et 20 ECU at session 8. Suppose that we randomly select sessions 3 and 6. We will pay you: (70 + 540) ECU x 0.1 €/ECU = 61€.

In each of the 8 market sessions, we will ask you to fill in a table indicating how many units of the asset you would like to buy or sell at each potential prices. As an indication, you will find at the bottom of the page, a table showing your final wealth in each of the two potential states of the economy (Heads or Tails) for various quantities traded. Recall that you can choose to trade any quantity between 0 and 4, even if this quantity is not indicated in the table.

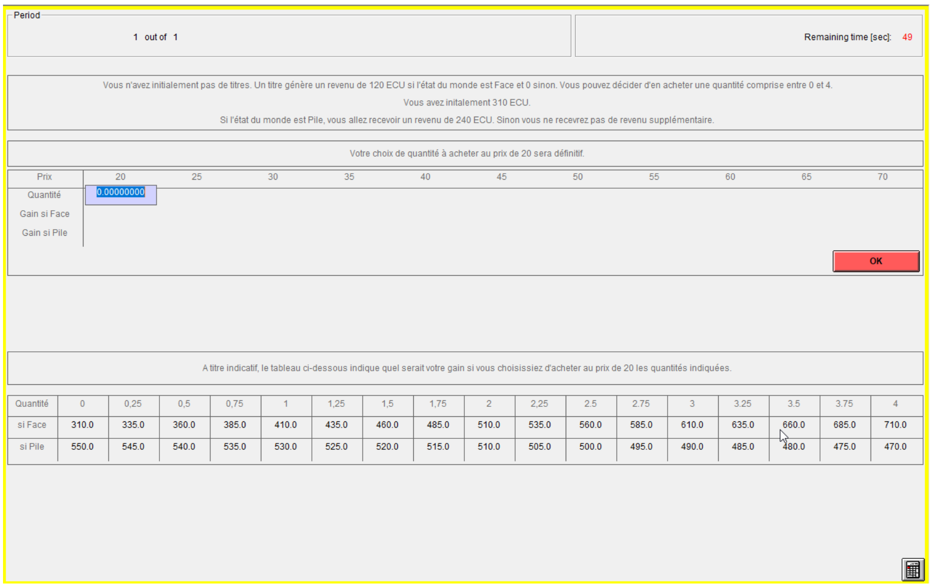
Period
1 out of 1
Remaining time [sec]: 49
Vous n'avez initialement pas de titres. Un titre génère un revenu de 120 ECU si l'état du monde est Face et 0 sinon. Vous pouvez décider d'en acheter une quantité comprise entre 0 et 4.
Vous avez initalement 310 ECU.
Si l'état du monde est Pile, vous allez recevoir un revenu de 240 ECU. Sinon vous ne recevrez pas de revenu supplémentaire.
Votre choix de quantité à acheter au prix de 20 sera définitif.
Prix 20 25 30 35 40 45 50 55 60 65 70
Quantité 0.00000000
Gain si Face
Gain si Pile
OK
A titre indicatif, le tableau ci-dessous indique quel serait votre gain si vous choisissiez d'acheter au prix de 20 les quantités indiquées.
Quantité 0 0,25 0,5 0,75 1 1,25 1,5 1,75 2 2,25 2.5 2.75 3 3.25 3.5 3.75 4
si Face 310.0 335.0 360.0 385.0 410.0 435.0 460.0 485.0 510.0 535.0 560.0 585.0 610.0 635.0 660.0 685.0 710.0
si Pile 550.0 545.0 540.0 535.0 530.0 525.0 520.0 515.0 510.0 505.0 500.0 495.0 490.0 485.0 480.0 475.0 470.0

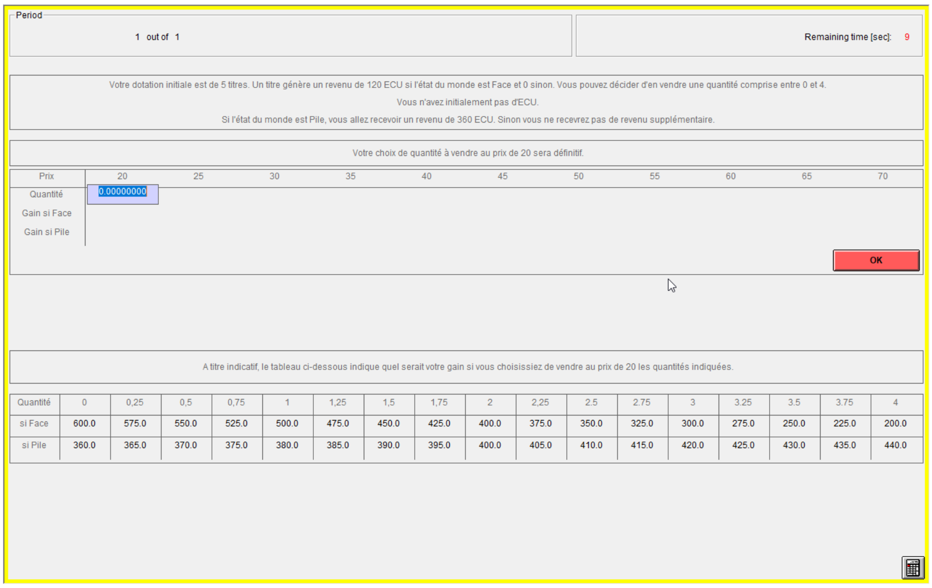
Period
1 out of 1
Remaining time [sec]: 9
Votre dotation initiale est de 5 titres. Un titre génère un revenu de 120 ECU si l'état du monde est Face et 0 sinon. Vous pouvez décider d'en vendre une quantité comprise entre 0 et 4.
Vous n'avez initialement pas d'ECU.
Si l'état du monde est Pile, vous allez recevoir un revenu de 360 ECU. Sinon vous ne recevrez pas de revenu supplémentaire.
Votre choix de quantité à vendre au prix de 20 sera définitif.
Prix 20 25 30 35 40 45 50 55 60 65 70
Quantité 0.00000000
Gain si Face
Gain si Pile
OK
A titre indicatif, le tableau ci-dessous indique quel serait votre gain si vous choisissiez de vendre au prix de 20 les quantités indiquées.
Quantité 0 0,25 0,5 0,75 1 1,25 1,5 1,75 2 2,25 2.5 2.75 3 3.25 3.5 3.75 4
si Face 600.0 575.0 550.0 525.0 500.0 475.0 450.0 425.0 400.0 375.0 350.0 325.0 300.0 275.0 250.0 225.0 200.0
si Pile 360.0 365.0 370.0 375.0 380.0 385.0 390.0 395.0 400.0 405.0 410.0 415.0 420.0 425.0 430.0 435.0 440.0

3

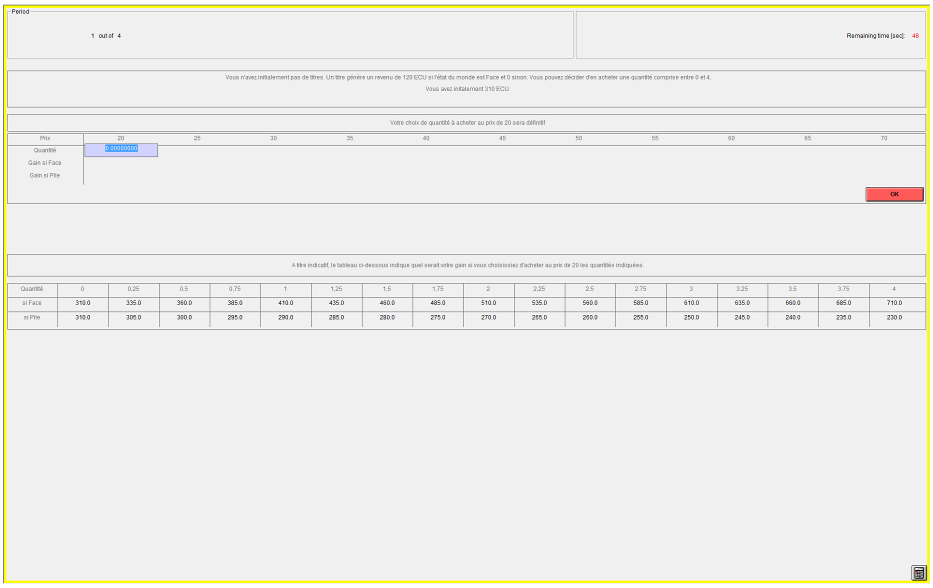

Period

1 out of 4

Remaining time [sec]: 48

Vous n'avez initialement pas de titres. Un titre génère un revenu de 120 ECU si l'état du monde est Face et 0 sinon. Vous pouvez décider d'en acheter une quantité comprise entre 0 et 4.

Vous avez initalement 310 ECU.

Votre choix de quantité à acheter au prix de 20 sera définitif

Prix	20	25	30	35	40	45	50	55	60	65	70
Quantité	0.00000000										
Gain si Face											
Gain si Pile											

OK

A titre indicatif, le tableau ci-dessous indique quel serait votre gain si vous choisissiez d'acheter au prix de 20 les quantités indiquées.

Quantité	0	0,25	0,5	0,75	1	1,25	1,5	1,75	2	2,25	2.5	2.75	3	3.25	3.5	3.75	4
si Face	310.0	335.0	360.0	385.0	410.0	435.0	460.0	485.0	510.0	535.0	560.0	585.0	610.0	635.0	660.0	685.0	710.0
si Pile	310.0	305.0	300.0	295.0	290.0	285.0	280.0	275.0	270.0	265.0	260.0	255.0	250.0	245.0	240.0	235.0	230.0

4

PRACTICE QUESTIONNAIRE

1. Consider a market session in which the transaction price is randomly determined. What is the likelihood that the transaction price equals 45?

a) 1 over 9 b) 1 over 10 c) 1 over 11 d) 1 over 12

What is the likelihood that the transaction price equals 70?

a) 1 over 9 b) 1 over 10 c) 1 over 11 d) 1 over 12

2. Suppose that we throw a coin three times and that we obtain Heads at each throw. What is the likelihood to obtain Tails at the next throw?

a) null b) 1 over 4 c) 1 over 2 d) 3 over 4

3. Suppose that we throw a coin four times and that we obtain Heads, Tails, Heads, and Tails. What is the likelihood to obtain Tails at the next throw?

a) null b) 1 over 4 c) 1 over 2 d) 3 over 4

4. Consider the table below that provides your final wealth depending on the quantity sold at the price of 55. Suppose that you propose to sell 1.25 at the price of 55. Suppose that the transaction price is 55, and that Tails comes up. What is your final wealth?

	0	0,25	0,5	0,75	1	1,25	1,5	1,75	2	2,25	2,5	2,75	3	3,25	3,5	3,75	4
FACE	610	594	578	561	545	529	513	496	480	464	448	431	415	399	383	366	350
PILE	370	384	398	411	425	439	453	466	480	494	508	521	535	549	563	576	590

N.B.: Face stands for Heads, and Pile stands for Tails.

OA.2 Robustness Analysis: The Experiment with a Coarse Price Grid

Table OA.1: **Experimental Protocol of the Experiment with a Coarse Price Grid**

Our additional experiment consists of five sessions and includes a total of 79 subjects who participated in eight replications. There are 11 prices on the coarse price grid $\{20, 25, \ldots, 65, 70\}$ with a tick size of 5. We alternated Treatment I in odd replications, consisting of Type 1 sellers and Type 2 buyers—plus possibly one Type 3 buyer when there is an odd number of subjects—and corresponding to no aggregate risk ("no AggR"), and Treatment II in even replications, consisting of Type 1 sellers and Type 3 buyers and corresponding to aggregate risk ("AggR"). For cohorts K and L, the first four replications involve a random mechanism, whereas the last four replications involve a call mechanism. For the other cohorts, the first four replications involve a call mechanism, whereas the last four replications involve a random mechanism. Column "# part." indicates the number of participants in each session.

		Round							
Session	# part.	1	2	3	4	5	6	7	8
I	17	Call				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
J	12	Call				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
K	15	Random				Call			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
L	20	Random				Call			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR
M	15	Call				Random			
		no AggR	AggR	no AggR	AggR	no AggR	AggR	no AggR	AggR

Table OA.2: **Distribution of Actions under the Coarse Price Grid and Associated Tests**

This table displays the distribution of actions under the coarse price grid. It is the equivalent of Table 4 corresponding to the baseline experiment. Panel A refers to actions at prices different from 60. The proportion of actions that are first-order stochastically dominated and nondominated is calculated. For Type 1 participants, for example, nondominated actions correspond to selling a quantity $q < 2$ at prices $S < 60$: in this quadrant, the bound is 0, nondominated actions far from $q = 2$ are $q \in]0, 1]$ and those that are near $q = 2$ are are $q \in]1, 2[$). The total number of actions is 760 in odd replications and 380 in even replication. Panel B corresponds to actions at a price of 60 at which there is no first-order stochastically dominated actions. The total number of actions is 76 in odd replications and 38 in even replication. In both panels, p-values are based on the binomial distribution with probability 0.5 under the null hypothesis. The number of trials correspond to the 8 replications, and the number of successes for H1 to the number of replications in which the data conforms with H1.

Panel A	Nondominated actions				Dominated actions		
(Price $S \neq 60$)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Replication	Bound	Far from $q = 2$	Near $q = 2$	$q = 2$	Near $q = 2$	Far from $q = 2$	Bound
1	117	142	114	129	87	129	42
2	61	91	69	78	45	25	11
3	115	143	159	197	68	58	20
4	60	103	98	71	32	13	3
5	120	134	174	196	77	41	18
6	65	76	103	79	29	22	6
7	137	125	162	199	78	37	22
8	68	60	98	77	43	28	6

Nonparametric tests Dominant vs dominated	Hypothesis H1	Successes H1	Trials	p-value under H0 (equally likely)
All	Prob [(1)+(2)+(3)] > Prob[(5)+(6)+(7)]	8	8	0.39%
Far from $q = 2$	Prob[(2)] > Prob[(6)]	8	8	0.39%
Near $q = 2$	Prob[(3)] > Prob[(5)]	8	8	0.39%
Bound	Prob[(1)] > Prob[(7)]	8	8	0.39%
Large vs Small	Prob[(5)] > Prob[(6)]	7	8	3.13%

Panel B	$q < 2$				$q > 2$		
(Price $S = 60$)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Replication	Bound	Far from $q = 2$	Near $q = 2$	$q = 2$	Near $q = 2$	Far from $q = 2$	Bound
1	13	18	10	11	11	9	4
2	1	3	4	11	6	9	4
3	6	7	9	23	16	12	3
4	2	0	10	11	8	7	0
5	4	10	12	19	20	7	4
6	2	5	5	12	5	8	1
7	3	6	12	28	13	9	5
8	1	2	8	11	8	5	3

Nonparametric tests $q < 2$ vs $q > 2$	Hypothesis H1	Successes H1	Trials	p-value under H0 (equally likely)
All	Prob[(1)+(2)+(3)] > Prob[(5)+(6)+(7)]	1	8	3.13%
Far from $q = 2$	Prob[(2)] > Prob[(6)]	2	8	10.94%
Near $q = 2$	Prob[(3)] > Prob[(5)]	1	8	3.31%
Bound	Prob[(1)] > Prob[(7)]	4	8	27.34%

Figure OA.1: **Aggregate Supply and Demand under the Coarse Price Grid**

Figure OA.1 illustrates the aggregate demand and supply schedules in our two treatments under the "Coarse" protocol. There were 38 Type 1 sellers and 38 Type 2 buyers in Treatment I with no aggregate risk, and 38 Type 1 sellers and 41 Type 3 buyers in Treatment II with aggregate risk. We average the quantity supplied by the sellers or demanded by the buyers at each price across the four replications of the same treatment.

Panel A: Treatment I (no aggregate risk)

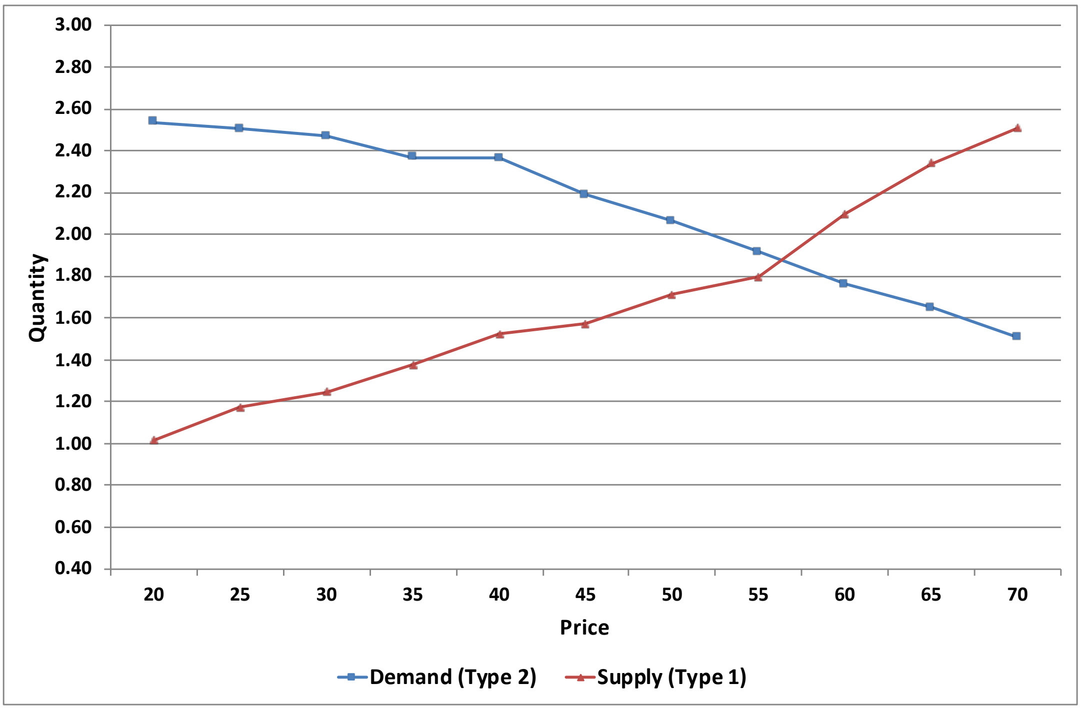

Panel B: Treatment II (aggregate risk)

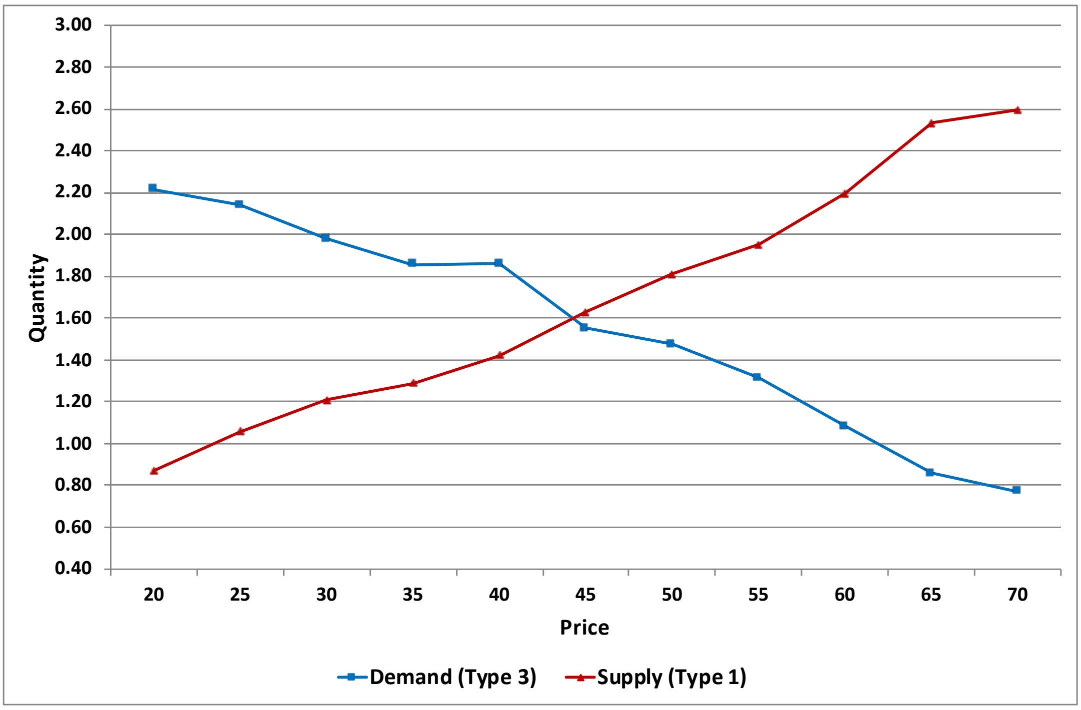

Figure OA.2: **Frequency of First-Order Stochastically Dominated Actions under the Coarse Price Grid**

Figure OA.2 illustrates the evolution of deviations from first-order stochastic nondominated actions by Type 1 sellers and Type 2 buyers in the additional sessions under the "Coarse" protocol. An action is first-order stochastically dominated when a Type 1 participant (respectively, a Type 2 participant) sells a quantity $q > 2$ at a price $S < 60$ or a quantity $q < 2$ at a price $S > 60$ (respectively, buys a quantity $q < 2$ at a price $S < 60$ or a quantity $q > 2$ at a price $S > 60$). For each subject and each replication, we compute the proportion of first-order stochastically dominated actions at each relevant price. Panels A and B show, for each replication, the average of this proportion across the 38 subjects of Types 1 and 2, respectively. Panel C shows the evolution of the proportion of small and large mistakes, where a small (respectively, a large) mistake is defined as a quantity supplied or demanded in a dominated quadrant that differs from 2 by less (respectively, more) than one unit. We use the standard deviation across all observations of each type to compute the 90% confidence intervals.

Panel A: Type 1 sellers (eight replications)

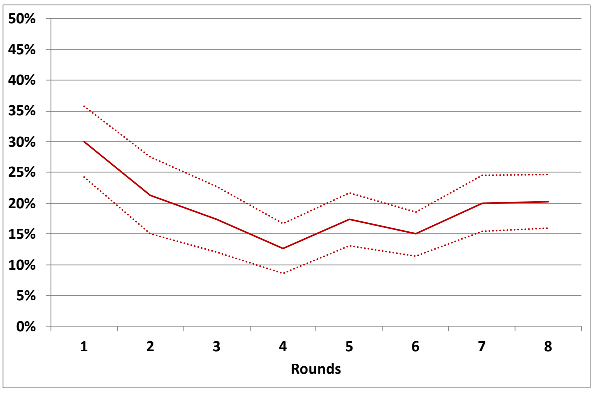

Panel B: Type 2 buyers (four replications)

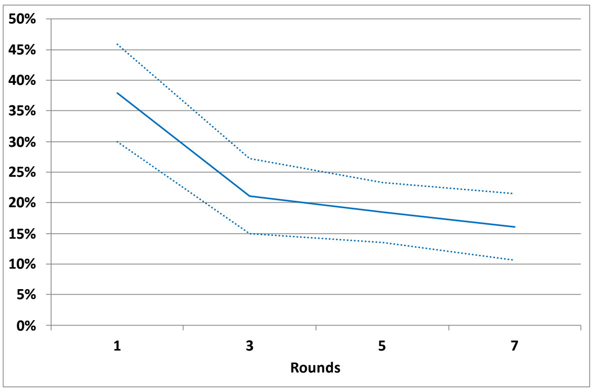

Panel C: Large and small mistakes for buyers and sellers