**Asset pricing and risk–sharing in a complete market:**

**An experimental investigation**

Bruno Biais, Thomas Mariotti, Sophie Moinas, Sébastien Pouget

Toulouse School of Economics

Very preliminary, comments welcome!

July 2014

Many thanks for helpful comments to seminar participants at City University, Zurich University and the 2014 Conference of the Financial Engineering & Banking Society in Guliford, Surry. Biais acknowledges financial support from the European Research Council under the European Community's Seventh Framework Program FP7/2007-2013 grant agreement N°295484.

**Abstract:**

We experimentally test the key implications of complete–markets competitive–equilibrium: Risks should be shared efficiently, individual wealth should comove with aggregate wealth and only aggregate risk should be priced. We elicit supply and demand curves from participants with heterogeneous random endowments and cross them to set (almost) market–clearing prices. When there is no aggregate risk, the market price of the risky asset tends towards its expected dividend, as predicted by theory. When there is aggregate risk, there is a risk premium of 8%, corresponding to an absolute risk–aversion coefficient of .55 for a power–utility representative agent. Participants, however, don't share risk fully efficiently.

# 1 Introduction

Since the seminal papers of Debreu (1959) and Arrow (1964), the general equilibrium theory of asset pricing and risk sharing in perfect and complete markets has offered an elegant framework and sharp implications: agents should share risk perfectly, which, in turn, implies only aggregate risk should be priced (for excellent textbook presentations, see, e.g., Huang and Litzenberger, 1988, or Demange and Laroque, 2006). Unfortunately, these implications are rejected by field data (for an insightful survey, see Kocherlakota, 1996). Is it because human cognition and preferences do not conform to the standard neoclassical micro–economic model? Or is it because, in practice, markets are strongly imperfect and incomplete? These two potential explanations have very different implications. The former calls for new models of human decision–making, deviating from expected utility maximization, while the latter calls for new models of imperfect and incomplete markets, populated by neo–classical agents.

While it is difficult to disentangle these two explanations with field data, in the lab it is possible to do so. In a controlled experimental setting one can make sure the market is perfect and complete. If, in that context, participants behaviour and market outcomes deviate from the predictions of neo–classical competitive equilibrium theory, the latter must be rejected. Otherwise, that theory must be pursued further, by incorporating incompleteness and imperfections. Evidence on this point would inform asset–pricing researchers as to which assumptions are relevant for theory and field–data analysis.

To take a first step in that direction, we conduct an experimental analysis of the simplest possible setting in which the basic tenets of the theory can be tested. There two final outcomes ($u$ and $d$) and two assets (the stock and the bond), so the market is complete. At the beginning of each of 8 trading rounds, participants start with heterogeneous endowments (stocks, bonds and other state–contingent

income). Individual demand and supply functions (specifying how many shares the participant wants to buy or sell at different prices) are elicited. Participants are asked to choose the quantity they want to buy or sell for at all prices on a grid. For half the rounds of the experiment, participants are told the market price will be randomly drawn and thus have no opportunity to manipulate the price.[1] During the other rounds, in contrast, participants are informed that the price is set to minimize the difference between aggregate supply and aggregate demand. Empirically, we observe similar behaviour within these two different price–setting mechanisms. This suggests participants are non–strategic, as in the standard competitive equilibrium model.[2]

We consider two treatments. In the first one, there is no aggregate risk: the sum of all endowments is constant across states. In the second treatment, there is aggregate risk: aggregate wealth is larger in state $u$ than in state $d$. The two treatments alternate from one round to the next. While market participants know their own endowment, they have no information about the endowments of the others or the alternance between the two treatments. This is in line with competitive Walrasian equilibrium, where agents only need to contingent decisions on the price, and don't engage in strategic interactions with one another.

Because the market is complete, when there is no aggregate uncertainty, risk–averse expected utility

[1] This is in the same spirit as the Becker, DeGroot, Marschak (1964). They i) ask a seller to quote a reservation price, ii) randomly draw a buying price, and iii) let the seller trades1 unit at the buying price if the latter is above the reservation selling price. While in that mechanism trades are only for one unit, our random price mechanism can be seen as an extension of the Becker, DeGroot, Marschak (1964) to general supply or demand curves.

[2] Our call market mechanism is similar to the market mechanism in Bohm et al (1997). They compare this market mechanism to the Becker, DeGroot and Marschak (1964) mechanism and find that the two yield similar transactions prices, except when participants are told they can place reservation selling prices that are clearly above any reasonable market price.

maximizing participants individual consumptions should hedge perfectly and their wealth should be constant across states. In this context, the price of the stock should be equal to the expectation of its final payoff, i.e., there should be no risk–premium. When there is aggregate uncertainty, in contrast, if participants are risk–averse there should be a positive risk–premium. Moreover, because the market is complete, individual consumptions should comove with aggregate consumption.

We ran the experiment with 69 students in March 2014, in Toulouse University. While individual behaviour is initially quite noisy, participants seem to converge towards the behaviour predicted by standard competitive equilibrium theory.

In the treatment with no aggregate risk, the distribution of individual quantities traded converges towards a mass at the full hedging trade. Thus, aggregate supply and demand converge towards the theoretical prediction, and so does the market price, which tends to the expected dividend.

In the treatment with aggregate risk, supply and demand curves are shifted in accordance with the predictions of the theory. This results in a market price significantly lower than the expected dividend, with a risk–premium of 8%, which, for power utility functions, corresponds to an average relative risk aversion coefficient of .55.

Yet, risk–sharing remains overall imperfect in both treatments, as some participants tend to trade more than predicted by risk–averse expected–utility maximization. In further work, we will investigate if this reflects that some participants are risk–loving, or if it stems from limited rationality.

Our work is directly in line with the insightful analyses of Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007). Like them we study the consequences of risk aversion in complete experimental markets and find prices involve risk–premium. Differences between the present paper and their Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007) include the following:

The trading protocol we consider is different from theirs. They consider a continuous double–auction, in which traders can dynamically post and hit quotes. This is similar to the workings of electronic limit order books during the day, and this is also in line with the seminal experimental studies of Smith Suchanek Williams (1988) and Plott and Sunder (1982, 1988). In contrast, we consider a call auction, in which participants submit supply and demand functions which are aggregated and crossed to determine the equilibrium price. This market mechanism is also used in practice, in general to open or close financial markets every day. In that centralized experimental market, to the extent that participants behave competitively (and our designs enables us to test precisely for that), competitive equilibrium precisely pins down individual behaviour and aggregate outcomes. This contrasts with the dynamic double auction, a bilateral market, which is in fact a complex game, the equilibrium of which has not been precisely characterized.

While Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007) consider a three–state model, we consider a two–state model. The advantage of a three–state complete market is that there are two risky assets, so that portfolios of risky assets can be analyzed. This enables Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007) to test (and reject) the hypothesis that agents hold the market portfolio. The problem of that design, however, is that participants must trade several risky assets. Such a market is experimentally difficult to organize and cognitively challenging for participants. Our design is simpler since we only have two assets: the risk–free bond and the risky stock. Thus agents need to trade in only one market, where the risky asset is exchanged for the risk–free bond. This enhances the simplicity and attractiveness of our call–auction design, where equilibrium can be precisely pinned down.

Bossaerts and Plott (2004), and Bossaerts, Plott and Zame (2007) offer very interesting evidence

on i) risk-premia and ii) failure of the two–fund separation theorem. The latter, however, obtains theoretically under specific assumptions about preferences, e.g., mean–variance. Rejection of the theoretical implication could therefore simply mean that agents' utility functions are not those assumed. In contrast, we test theoretical implications that hold for all utility functions as long as agents are expected utility maximizers.

The present paper is also related to the experimental literature studying in the lab the consequences of risk–aversion for economic and financial decisions, e.g., Holt and Laury (2002). Economists sometimes raise doubts about the significance of risk–aversion in the lab, where stakes are relatively low. To shed light on this, Holt and Laury (2002) elicit risk aversion by questionnaires, sometimes involving high stakes and sometimes low stakes.[3] They find that, even at low payoff levels (when prizes are below \$4) two thirds of the participants exhibit risk–aversion. Assuming power utility, the majority of the participants have estimated relative risk–aversion between .15 and .97 (both for the low payoff treatment and for the treatment where payoffs are scaled up by a factor 20). Our estimates are of the same order of magnitude as theirs.[4]

Section 2 presents our theoretical and experimental framework. Section 3 presents our experimental findings. Section 4 briefly concludes.

[3] Bossaerts and Plott (2004) also find that participants exhibit risk aversion both for relatively low stakes and for much larger stakes.

[4] This is in line with other estimates of relative risk aversion in experiments around .5. For private value auctions see Cox and Oxaca (1996) and Goeree et al (2002), for asymmetric matching penny games see Goeree et al (2003), and for one–shot matrix games see Goeree and Holt (2004).

# 2 Theoretical and Experimental Framework

## 2.1 Theoretical Background

Consider a one-period complete-market pure-exchange economy populated by competitive risk-averse agents maximizing their expected utility over final wealth, given common beliefs about finitely many states of nature. In this context, a competitive equilibrium exists and is Pareto efficient. That is, risks are efficiently shared in equilibrium.

A striking implication of efficient risk sharing is that individual state-contingent wealth co-moves with aggregate wealth. This reflects that the rankings of individual wealth across states of nature must be the same for all agents, which in turn reflects that the rankings of marginal utilities must be the same. If they were not, marginal rates of substitution between states of nature would not be equalized across agents, which would contradict efficiency.

This co-monotonicity property has sharp asset-pricing implications. If aggregate wealth is constant across states of nature, even if individual endowments may vary with the state of nature, individual consumption is constant across states. Since agents do not bear risk, the equilibrium price of any asset is equal to its expected final payoff. In contrast, if aggregate wealth is not constant across states of nature, assets typically command a nonzero risk premium. For instance, asset whose final payoff co-varies positively with marginal utilities of final wealth are priced below their expected final payoff.

## 2.2 Experimental Design

We design a simple experimental market to test these predictions, keeping as close as possible to the competitive-equilibrium framework.

**Assets** Markets are complete. There are two equally likely states of nature, $\omega = u, d$, and two assets, a bond and a stock. 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 one market, in which the stock is traded against the bond; that is, we take the bond as the numéraire.[5] We denote by $S$ the price of the stock relative to the bond.

Although considering only two assets precludes us from analyzing portfolios of risky assets, it sidesteps the difficulties associated with simultaneous trading in several experimental markets that have been emphasized by Bossaerts and Plott (2004) or Bossaerts, Plott, and Zame (2007). Moreover, having only one market in which the stock can be traded against the bond simplifies the cognitive task faced by participants. If, by contrast, we had considered two markets in which the stock and the bond could be traded against a perishable unit of account, participants would have faced a more difficult task: they would have had to factor their trades on the two markets into their budgets constraints, and to compare the prices on these two markets.

**Endowments** Participants' endowments come from their initial holdings of bond and stock, and from state-contingent additional income. We will consider two treatments. In Treatment 1, there is no aggregate risk: the sum of individual endowments is constant across states of nature. In Treatment 2, there is aggregate risk: the sum of individual endowments is larger in state $u$ than in state $d$. We chose this design because, as explained in the next subsection, it allows us to test sharp predictions of the theory.

For simplicity, participants can receive only three types of endowments:

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

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 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 $d$.[6]

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

In Treatment 1, 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.

In Treatment 2, if the number of participants is even, there are Type 1 and Type 3 participants, in equal numbers; whereas if the number of participants is odd, there is one additional Type 3 participant.

Participants do not know that there are two treatments. At each replication of the experiment, they are only informed about their own endowments and the distribution of the dividend.

**Supply and Demand Curves** To stay as close as possible to the competitive-equilibrium model and to gather precise data on individual behavior, we elicit supply and demand curves. As explained below, if all agents are competitive and risk averse, in equilibrium, Type 1 agents never buy the stock, while Type 2 and Type 3 agents never sell it. To simplify the task of participants in the experiment, we therefore restrict Type 1 participants to supply, and Type 2 and 3 participants to demand, nonnegative quantities of the stock. A further simplification is that participants are restricted to trade quantities

[6] The endowments for Type 1 and Type 2 participants have been chosen such that the traded quantity at which they do not have any risk is equal to 2. See below for more details.

no greater than 4.[7] In the experiment, participants are asked which quantity they want to sell or buy at each point of a price grid ranging from 52 to 62 with unit increments. Whereas the price grid is discrete, at each price $S$, we allow the 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 in principle to equate marginal utility and price when they select trades in $[0, 4]$.

**Pricing** The closest design to the competitive-equilibrium model would be to aggregate supply and demand curves, and set the price such that the market clears. A first difficulty is that, in our experiment, the price grid is discrete, so that exact market clearing may fail to obtain. A second difficulty is that participants may behave strategically and try to manipulate the price. To address these issues, we consider two different pricing mechanisms.

1. In the *call* mechanism, the price is set to minimize the gap between supply and demand.[8] At this price, there is typically rationing. To avoid this consideration to affect the choices of the participants, we supply or demand the quantity of stock needed to clear the market. Thus participant orders are fully executed.

2. In the *random* mechanism, the price is drawn from the price grid, each price being equally likely. As in the call mechanism, participant orders are fully executed.

[7] These restrictions, in particular, forbid short sales, which would anyhow not arise in equilibrium. We cap trades at 4, whereas Type 1 participants initially hold 5 shares of the stock, so that, in Treatment 1, the equilibrium price is the same when agents use quantal responses as when they perfectly maximize expected utility. This feature enhances the robustness of theoretical predictions.

[8] If there is more than one price at which the gap between supply and demand is minimized, we choose the price that yields the maximum volume. If there is still an indeterminacy, we randomly pick one of the remaining candidates.

At the beginning of each replication of the experiment, all participants were told which mechanism will be used to set the price. Strategic considerations could affect participants' behavior when they knew their orders will be processed in the call mechanism, but not when they knew the random mechanism will be used to set the price. Therefore, comparing the two enables us to test whether strategic considerations significantly affected participants' behavior. If they did not, then the call mechanism provides a good approximation of a competitive market.

## 2.3 Testable Implications

### 2.3.1 Individual Behavior

Our experimental design involves two equally probable states. This allows us to express the objective function of an expected-utility maximizer as a function of the expectation and variance of his final wealth. Denote by $W^i(u)$ and $W^i(d)$ the final wealth in states $u$ and $d$ of agent $i$, and let

$$\mu^i \equiv \frac{1}{2}\left[W^i(u) + W^i(d)\right] \tag{1}$$

and

$$\sigma^i \equiv \frac{1}{2}\left|W^i(u) - W^i(d)\right| \tag{2}$$

be the expectation and standard deviation of final wealth. Then the expected utility of agent $i$ with utility function $v^i$ writes as:

$$\mathbf{E}[v^i(W^i(\omega))] = \frac{1}{2}\left[v^i(\mu^i + \sigma^i) + v^i(\mu^i - \sigma^i)\right]. \tag{3}$$

Therefore, our symmetric two-state setup allows for a mean-variance representation without requesting parametric assumptions on the utility functions $v^i$.

**Type 1** If a Type 1 participant sells $Q$ shares of the stock at price $S$, his final wealth is

$$W^i(u) = 600 + (S - 120) \times Q, \tag{4}$$

$$W^i(d) = 360 + S \times Q, \tag{5}$$

so that

$$W^i(u) - W^i(d) = 120 \times (2 - Q).$$

Type 1 participants are perfectly hedged if they sell 2 shares. When $S = 60$, it is optimal for a risk-averse expected-utility maximizer Type 1 participant to do so, thus obtaining full insurance at an actuarially fair price. What if the stock price is not equal to the expected dividend? According to (1)–(2), $\mu^i = (S - 60) \times Q + 480$ and $\sigma^i = 60 \times |Q - 2|$. Therefore, two possible trades $Q^+ = 2 + \varepsilon$ and $Q^- = 2 - \varepsilon$ for $\varepsilon \in (0, 2]$ lead to equally volatile final wealth. However, for $S < 60$, trading $Q^+$ shares leads to a lower expected final wealth than trading $Q^-$ shares, while for $S > 60$, the converse holds. Therefore, by (3), as soon as an agent is maximizing expected utility, trading $Q^+$ shares is dominated by trading $Q^-$ shares when $S < 60$, while the converse holds when $S > 60$. This yields the following implication

**Implication 1:** *If a Type 1 participant is an expected-utility maximizer, he does not supply more than* 2 *shares at* $S < 60$, *nor supply less than* 2 *shares at* $S > 60$.

Note that Implication 1 holds no matter the shape of the utility function $v^i$, as long as it is increasing in final wealth. Therefore, it is equally valid for risk-averse and risk-loving participants, or for participants who are risk-averse over some ranges of wealth, and risk-loving on others.

To illustrate this implication, it is useful to represent the choices of Type 1 agents in the price–quantity plane, as done in Figure 1. The relevant prices are between 52 and 62, and the relevant

quantities are between 0 and 4, and there are four quadrants determined by the horizontal line $Q = 2$ and the vertical line $S = 60$. If follows from Implication 1 that the North-West and South-East quadrants are dominated for an expected-utility maximizer. A risk-averse Type 1 agent's supply function goes through the point $(60, 2)$. A risk-loving Type 1 agent's supply of the stock equals 0 for $S < 60$ and 4 for $S > 60$; at $S = 60$, he is indifferent between $Q = 0$ and $Q = 4$.[9]

**Type 2** If a Type 2 participant buys $Q$ shares of the stock at price $S$, his final wealth is

$$W^i(u) = 310 + (120 - S) \times Q, \tag{6}$$

$$W^i(d) = 550 - S \times Q, \tag{7}$$

so that

$$W^i(u) - W^i(d) = 120 \times (Q - 2).$$

Exploiting the symmetry between (4)–(5) and (6)–(7), we obtain the following implication.

**Implication 2:** *If a Type 2 participant is an expected-utility maximizer, he does not demand less than* 2 *shares at* $S < 60$, *nor demand more than* 2 *shares at* $S > 60$.

This implication is illustrated in Figure 2, which is the mirror image of Figure 1.

**Type 3** If a Type 3 participant buys $Q$ shares of the stock at price $S$, his final wealth is

$$W^i(u) = 310 + (120 - S) \times Q, \tag{8}$$

$$W^i(d) = 310 - S \times Q, \tag{9}$$

[9] This last observation stems from the symmetry of our quantity range around $Q = 2$.

so that

$$W^i(u) - W^i(d) = 120 \times Q.$$

In contrast with Type 1 and Type 2 participants, Type 3 participants start with a riskless initial endowment. Thus, as shown by (8) and (9), buying stocks increases their exposure to risk. The restrictions imposed by theory are much weaker than for Type 1 and Type 2 participants. If a Type 3 participant is a risk-averse expected-utility maximizer, his demand of stock at $S \geq 60$ is 0, whereas if a Type 3 participant is a risk-loving expected-utility maximizer, his demand of stock at $S \leq 60$ is 4. These implications are illustrated in Figure 3.

#### 2.3.2 Equilibrium Outcomes

**Treatment 1** In Treatment 1, there is an equal number of Type 1 and Type 2 participants, and possibly a Type 3 participant. Aggregate wealth is thus constant across states. Competitive equilibrium with risk-averse agents yields the following testable implication.

**Implication 3** *In Treatment 1, if participants are risk-averse expected-utility maximizers, there exists an equilibrium such that $S = 60$ and individual final wealth is equalized across states of nature, that is, $W^i(u) = W^i(d)$ for all $i$. In this equilibrium, Type 1 participants sell 2 shares, Type 2 participants buy 2 shares, and Type 3 participants do not trade.*

If all participants are finitely and strictly risk-averse expected-utility maximizers, the equilibrium described in Implication 3 is the unique equilibrium: this is an instance of the mutuality principle (Borch (1962)). The intuition is that because aggregate wealth is constant across states and markets are complete, in equilibrium agents fully diversify their risk and, therefore, there is no risk premium

in the stock price. Whereas Implication 3 obtains 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 pricing grid.

Uniqueness of equilibrium is not satisfied in some extreme cases. For instance, if all participants are infinitely risk averse, their supply and demand of stock are perfect inelastic (at $Q = 2$ for Type 1 and Type 2 participants, and at $Q = 0$ for Type 3 participants). In that case, the equilibrium price is indeterminate. By contrast, if all participants are risk neutral, then the equilibrium price of the stock has to be equal to its expected dividend, $S = 60$, but at this price, supply and demand are perfectly elastic, and the equilibrium trading volume is indeterminate.

**Treatment 2** In Treatment 2, there are only Type 1 and Type 3 participants. Aggregate wealth is strictly greater in state $u$ than in state $d$. Competitive equilibrium with risk-averse agents yields the following testable implication.

**Implication 4** *In Treatment 2, if participants are risk-averse expected-utility maximizers, there exists an equilibrium such that* $S \leq 60$ *and individual final wealth is co-monotonic with aggregate wealth, that is,* $W^i(u) \geq W^i(d)$ *for all* $i$.

If all participants are finitely and strictly risk-averse expected-utility maximizers, then $S < 60$ and $W^i(u) > W^i(d)$ for all $i$. The gap $60-S$ between the expected dividend of the stock and its equilibrium price is the risk premium requested by risk-averse agents to bear aggregate risk. Uniqueness of equilibrium is not guaranteed in general, but if relative risk aversion $-wv^{i\prime\prime}(w)/v^{i\prime}(w)$ is less than 1 for all $w$ and $i$, then substitution effects dominate income effects and the equilibrium is unique.[10]

[10] Technically, this condition ensures that the aggregate aggregate excess demand function for state-contingent wealth

## 3 Experimental findings

We ran the experiment on March 12, 2014. The participants were 69 students enrolled in the first year of Toulouse University's master in finance.[11] There were four groups, each with 15 to 20 participants. Each group participated in 8 replications of the experiment, lasting, overall, one hour and a half. Treatment 1 and Treatment 2 alternated during the 8 replications. For two of the four groups, the first four replications involved random prices, while the last four replications involved a call auction. For the two other groups, the first four replications involved a call auction, while the last four replications involved random prices.

For each group, two of the eight rounds were randomly drawn, and participants received the sum of their final wealths in these two rounds divided by ten. The average individual payment was € 87, the minimum was € 53 and the maximum € 115. We ran an anonymous survey among the participants, asking them what their weekly budget was. The minimum was € 65, the median € 475, the mean € 483 and the maximum €1000. Thus, the amount participants could make in the lab, and its variability, were significant relative to their budget in the field.[12]

satisfies the gross substitute property, see Varian (1985) or Mas-Colell, Whinston, and Green (1995, Example 17.F.2 and Proposition 17.F.3, pp. 612–613).

[11] These students come from different backgrounds. The majority studied management, and had very little exposure to micro-economic or finance theory. Some studied economics, and had greater exposure to micro. Others come from engineering or maths.

[12] Bossaerts and Plott (2004) compare participants' behaviour in asset market in the US and in Bulgaria. The experimental design was the same in the two countries, but monetary incentives were much stronger in Bulgaria. Bossaerts and Plott (2004) find qualitatively similar results. In both countries there is a risk–premium. The main difference is that the risk–premium is larger in Bulgaria.

### 3.1 Individual supply and demand functions

There is a lot of heterogeneity in individual behavior. To illustrate this, Figure 4 plots the supply functions of four different participants in four replications of Treatment 1. The behavior of Participant $a$ seems quite noisy, and involves several dominated actions. In contrast, the behavior of Participant $b$ is very steady and conforms to the prediction of theory for an infinitely risk–averse expected–utility maximizer. Yet another contrasting example is offered by Participant $c$, whose behavior conforms to the prediction of theory for a risk–loving agent. There are also many participants whose behavior is in line with the predictions of theory for finitely risk–averse expected utility maximizers, for example Participant $d$.

In 16 replications of the experiment, the price at which participants traded was randomly drawn. For the other 16 replications the price was set to minimize the gap between supply and demand. To test the null hypothesis of competitive behavior, we ran Wilcoxon rank–sum tests comparing the distribution of quantities offered or demanded in the two price setting mechanisms.[13] As illustrated in Figure 5, the $p$–value was above 10% in all cases but 2. Therefore one cannot reject the null hypothesis of competitive behavior. Accordingly, when analyzing aggregate outcomes we hereafter pool the observations from the two pricing mechanisms.

To document average behavior across individual participants, Figure 6 reports the percentage of

[13] At each price level, we first computed for each individual the average quantity in each treatment (aggregate risk vs no aggregate risk) and price setting mechanism (call vs random). So we computed an average quantity across 2 replications for each subject. Then, for each price level, we ran 4 separate Wilcoxon rank–sum tests: buyer/aggregate risk, buyer/no aggregate risk, seller/aggregate risk and seller/no aggregate risk. In each of these tests we compared the distribution across subjects of quantities for the random price mechanisms to its counterpart for the call mechanism. The number of observations (individuals) for the 4 tests ranged between 32 and 36.

individual actions that are dominated for Type 1 and Type 2 participants.[14] Note that for Type 3 participants there are no dominated actions, we only report evidence for the treatment with no aggregate risk. The figure shows that while the percentage of dominated action is initially quite large (45% for Type 1 and 31% for type 2), it declines with experience (to 23% for Type 1 and 15% for Type 2). This suggests participants are learning not to play dominated actions, but not perfectly.

## 3.2 Market prices

Figure 7 and Figure 8 document market pricing. For each of the two treatments (Treatment 1 without aggregate risk, and Treatment 2 with aggregate risk), each of the four groups played four replications.

In Treatment 1, the prediction from theory is quite sharp: as soon as participants maximize expected utility, the market price must be equal to the expected dividend, 60. Figure 7 depicts the prices minimizing the gap between supply and demand for the 16 replications of Treatment 1.[15] At the first replication, prices were quite low, ranging between 52 and 58. But, as participants got more experienced, prices tended to increase. At the fourth replication of Treatment 1, market prices ranged between 58 and 62. This suggests the market converged towards the theoretical equilibrium price as participants got more experienced. This is consistent with individual learning documented above.

For Treatment 2, the prediction from theory is less sharp. The only thing one can say is that, if participants are risk–averse expected–utility maximizers, then the market price should be lower than

[14] For Type 1 (resp. Type 2), there is a dominated action if the participant supplies (resp. demands) strictly more (resp. less) than 2 for $S < 60$, or strictly less (resp. more) than 2 for $S > 60$. For each participant and each period, we count the number of dominated actions across the 10 relevant prices. Then for each period, we average the number of dominated actions across participants. Finally we convert this number into a percentage.

[15] Each of the four groups of participants played four replications.

60. This is what we observe in Figure 8, which depicts the prices minimizing the gap between supply and demand for the 16 replications of Treatment 2.

To shed light on aggregate demand and supply in our experimental market, we aggregate all the supply and demand curves from all Type–1 and Type–2 participants in Treatment 1, and from all Type–1 and Type–3 participants in Treatment 2. Figure 9 depicts aggregate demand and supply in Treatment 1 (divided by the number of participants–replications, to facilitate interpretations). Demand is approximately decreasing and supply approximately increasing. Quite strikingly, the price minimizing the gap between supply and demand is 60, which is the equilibrium price predicted by theory. At that price, trading volume per participant is slightly larger than what equilibrium would entail (two shares).

Figure 10 depicts aggregate demand and supply per participant in Treatment 2. As predicted by theory, supply is not very different from its counterpart in Figure 9. In contrast, as predicted by theory, demand is much lower. Correspondingly, the price minimizing the gap between supply and demand (55) is lower than the expected dividend. This is a risk–premium of 8%. For a representative investor endowed with aggregate wealth and with power utility function, this corresponds to a relative risk aversion coefficient of .55. This is quite in line with evidence from other experiments infering risk–aversion from participants' behavior.[16]

[16] See Holt and Laury (2002) for lottery choices, Cox and Oxaca (1996) and Goeree et al (2002) for private value auctions , Goeree et al (2003) for asymmetric matching penny games, and Goeree and Holt (2004) for one–shot matrix games.

### 3.3 Transactions at the aggregate market price

Next, we investigate to which extent, at the aggregate (almost) market–clearing prices discussed just above, participants' trades are as implied by theory.

For Treatment 1, Figure 11 depicts the distribution of individual Type–1 supplies and Type–2 demands at $S = 60$. Theory predicts all Type–1 and Type–2 participants should trade two shares, and thus perfectly hedge at that actuarially fair price. As can be seen in Figure 11, the transaction sizes most frequently observed are between 1.75 and 2.25, in line with theory. There is, however, quite a bit of dispersion around this theoretically predicted outcome. What's more, the frequency of trades above 2.25 is larger than the frequency of trades below 1.75. This suggests that, to some extent, there is excess trading (in line with our above remark about transaction size in Figure 9.)

Figure 12 depicts the distribution of participants' trades in Treatment 2. In that treatment, the implication of theory is less sharp than in Treatment 1. Note however that the frequency with which participants buy less than 2 shares is larger than in Treatment 1, which is in line with theory.

Figure 13 documents learning in Treatment 1. Panel A shows the distribution of transaction sizes in the first two replications, while Panel B shows its counterpart for the last two replications. The frequency of trades close to 2 is higher, and the frequency of excessively large trades is lower, in Panel B than in Panel A. This suggests participants learn to converge towards optimal behaviour. Figure 14 shows a similar pattern for sellers in Treatment 2.

## 4 Conclusion

Our experimental design offers the simplest possible setting to test the key implications of complete–markets, competitive–equilibrium theory. Crossing aggregate supply and demand curves, we get prices that are, overall, consistent with the implications of theory. Yet, while participants share some risk, they leave some risk–sharing opportunities unexploited, in contrast with theory. In further work, we will examine whether this reflects risk–loving or deviations from expected utility maximization.

## 5 References


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

Becker, G., M. DeGroot and J. Marschak (1964). "Measuring utility by a single–response sequential method." *Behavioral Science.* 9, 226–32.

Bohm, P., J. Linden and J. Sonnegaerd (1997). "Eliciting reservation prices: Becker–DeGroot–Marschak mechanisms versus markets." *The Economic Journal*, 107, 1079–1089.

Borch, K., (1962). "Equilibrium in a reinsurance market", *Econometrica*, 30 (3) 424–444.

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

Bossaerts, P., C. Plott, and W.R. Zame (2007): "Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments," *Econometrica*, 75(4), 993–1038.

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

Cox, J., and R. Oaxaca (1996). "Is bidding behavior consistent with bidding theory for private

value auctions?" in R. Isaac, ed., *Research in Experimental Economics*, 6, 131–48.

Debreu, G. (1959). *The Theory of Value: An axiomatic analysis of economic equilibrium.*

Demange, G., and G. Laroque (2006). *Finance and the Economics of Uncertainty*, Wiley-Blackwell.

Goeree, J. and C. Holt (2004). "A Model of Noisy Introspection." *Games and Economic Behavior*, 46, 365-382.

Goeree, J., C. Holt and T. Palfrey (2003). "Risk Averse Behavior in Generalized Matching Pennies Games." *Games and Economic Behavior*, 45, 97-113.

Goeree, J., C. Holt and T. Palfrey (2002). "Quantal Response Equilibrium and Overbidding in Private-Value Auctions" *Journal of Economic Theory*, 104, 247-272.

Holt, C. and S. Laury (2002). "Risk–aversion and incentive effects." *The American Economic Review*, 92, 1644–1655.

Huang, C.F., and R. Litzenberger (1988). *Foundations for financial economics*, North–Holland.

Kocherlakota, N. (1996). "The equity premium: It's still a puzzle." *Journal of Economic Literature*, 34, 42–71.

Mas-Colell, A., M.D. Whinston, and J.R. Green (1995): *Microeconomic Theory*, New York, Oxford: Oxford University Press.

Plott, C. and S. Sunder (1982). "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *Journal of Political Economy,* 90, 663-698.

Plott, C. and S. Sunder (1988). "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets." *Econometrica*, 56, 1085-1118.

Smith, V, G. Suchanek and W. Williams (1988). "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica*, 56, 1119-1151

Varian, H. (1985): "Additive Utility and Gross Substitutes," Unpublished Manuscript, University of Michigan, Ann Arbor.

## Appendix

Consider a one-period complete-market environment with two states of the world, $\omega = u, d$, and two assets, a stock paying dividend $D(\omega)$ in state $\omega$, with $D(u) > D(d)$, 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 additional income $I(\omega)$ in state $\omega$. In the market, an agent can trade to his 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

$$p_s(\theta_{1,s} - \theta_{0,s}) + p_b(\theta_{1,b} - \theta_{0,b}) \leq 0. \tag{10}$$

The agent maximizes his utility from state-contingent final wealth

$$W(\omega) \equiv I(\omega) + \theta_{1,s} D(\omega) + \theta_{1,b}, \quad \omega = u, d. \tag{11}$$

Because utility is increasing in final wealth in each state, the budget constraint (10) is binding. Hence

$$\theta_{1,b} = \theta_{0,b} - \frac{p_s}{p_b}(\theta_{1,s} - \theta_{0,s}). \tag{12}$$

Substituting (12) into (11), we get

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

This shows that the only factor that affects the final wealth of the agent in state $\omega$, beyond his initial endowment $I(\omega) + \theta_{0,b} + \theta_{0,s} D(\omega)$, is the product of his 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, $D(\omega) - p_s/p_b$. Therefore, because there are only two states of the world, the simple market structure in our experimental setting, where participants can only trade the stock against the bond, implements a complete market structure. Letting $S \equiv p_s/p_b$ and $Q \equiv |\theta_{1,s} - \theta_{0,s}|$, applying (11) to Type 1, Type 2, and Type 3 participants yields (4)–(5), (6)–(7), and (8)–(9).

# Figure 1: Type 1

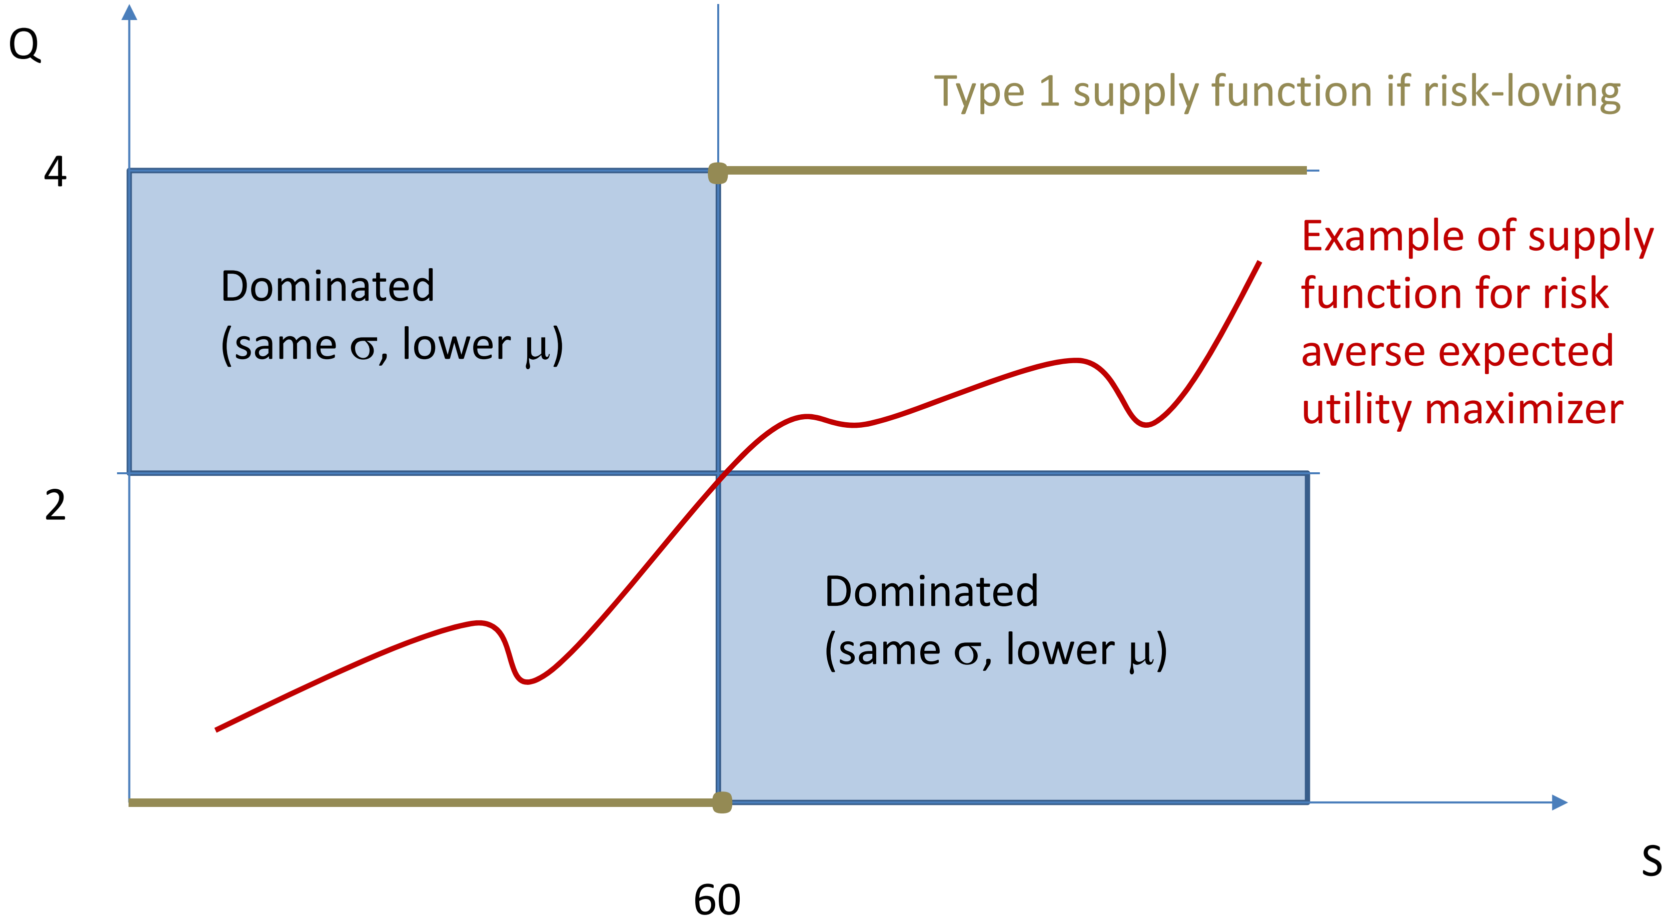

# Figure 2: Type 2

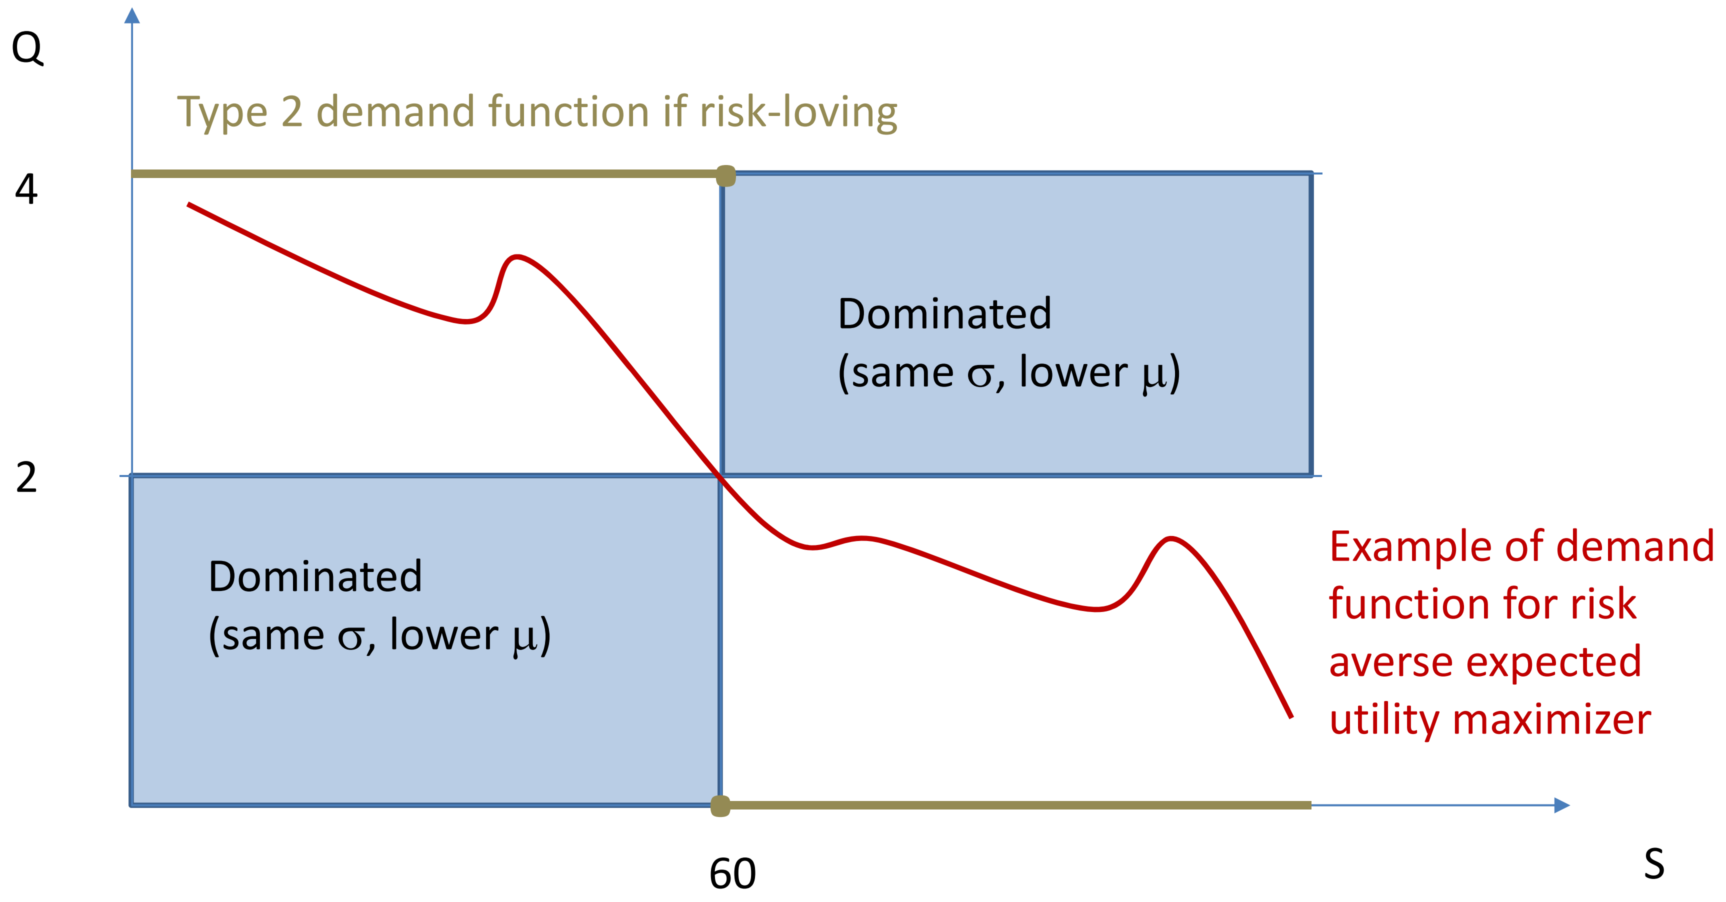

# Figure 3: Type 3

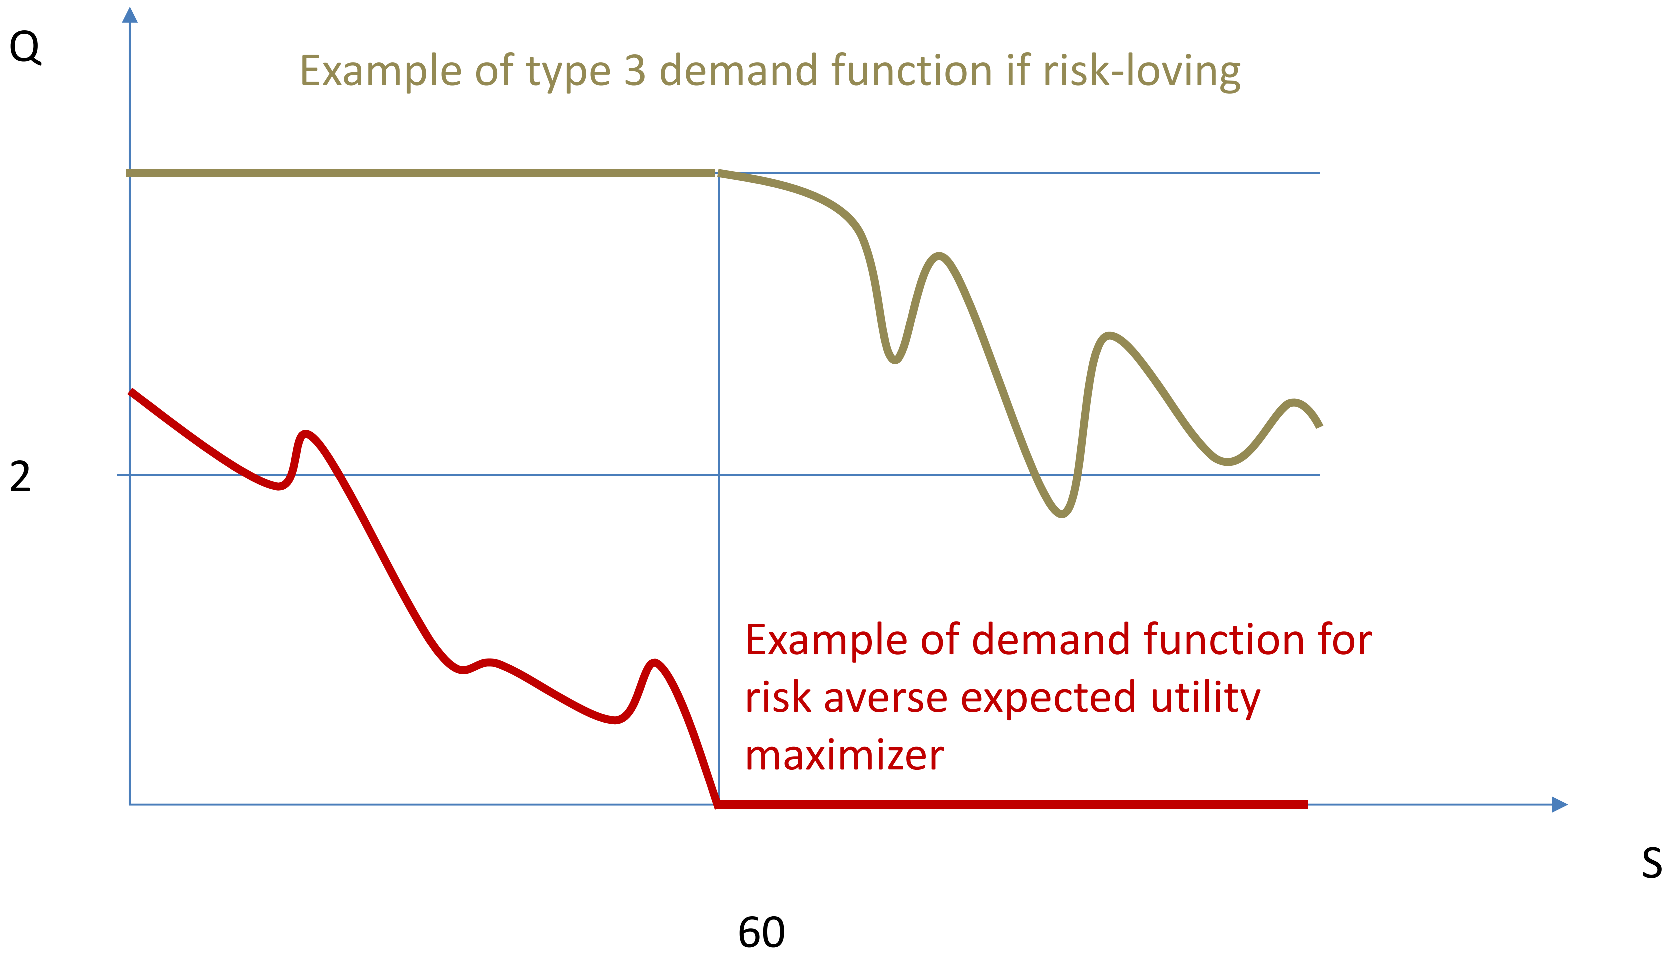

# Figure 4: 4 examples of Type 1 participants choices

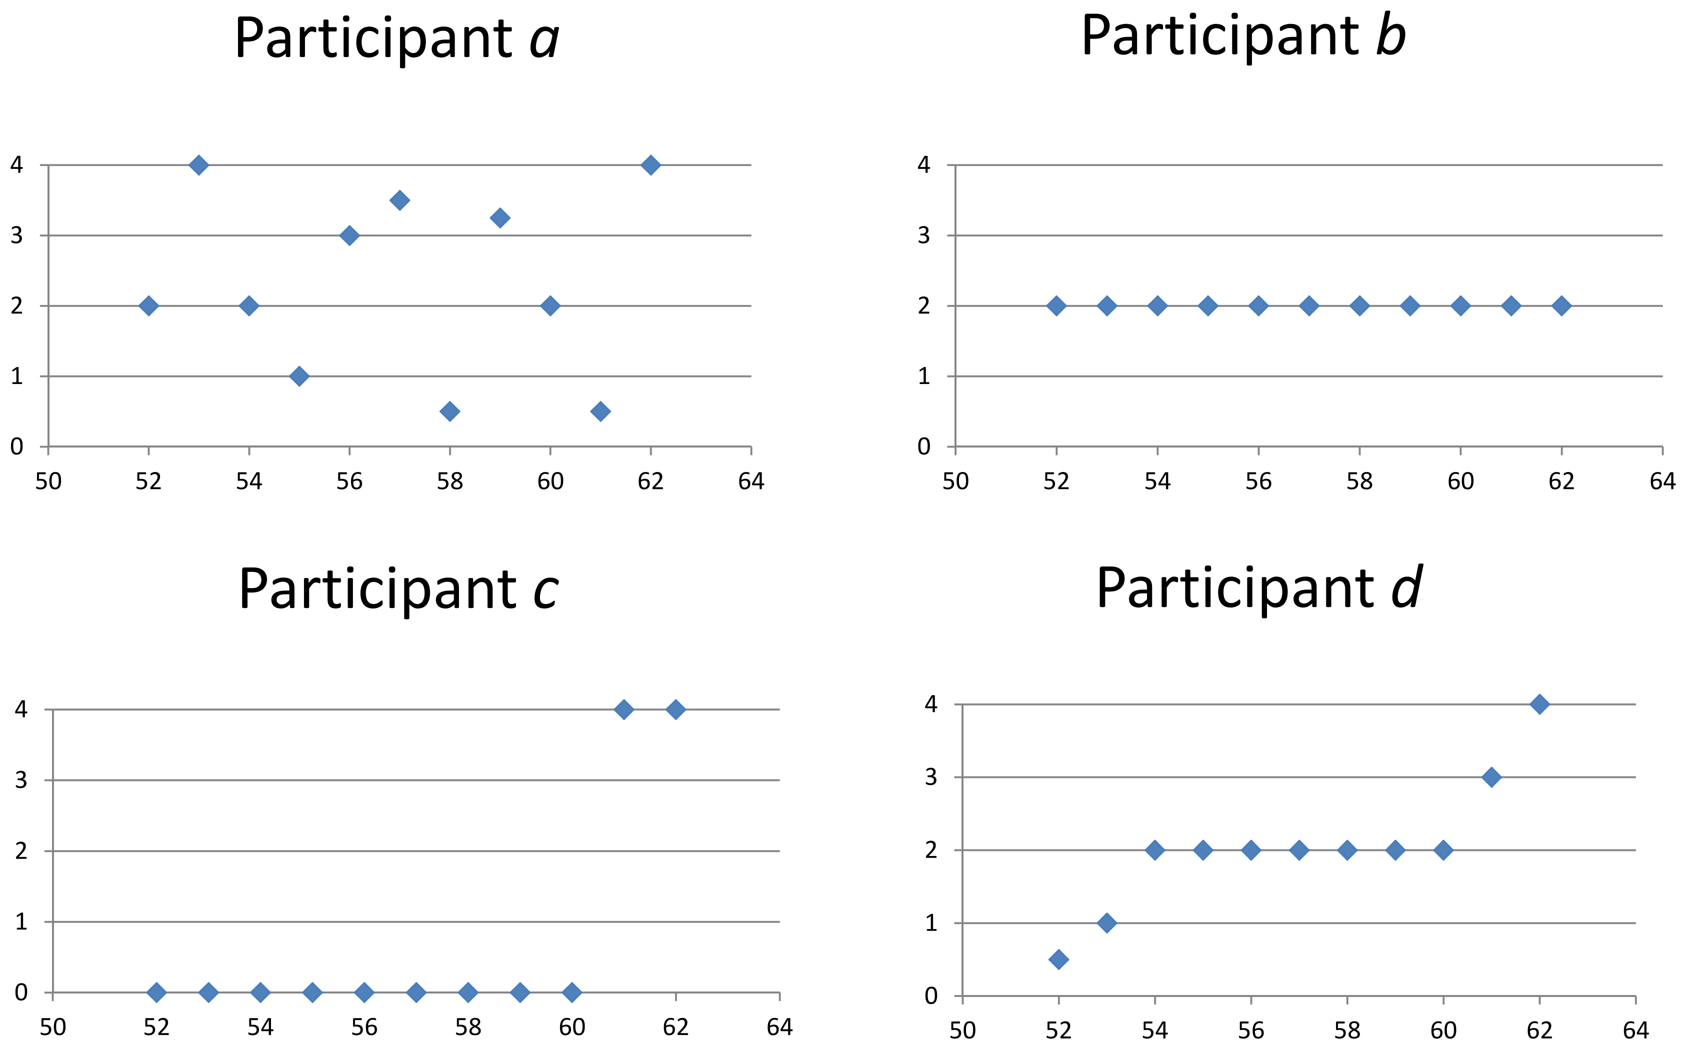

## Figure 5: p-values for Wilcoxon rank-sum test of H0 that call and random price mechanisms are not significantly different

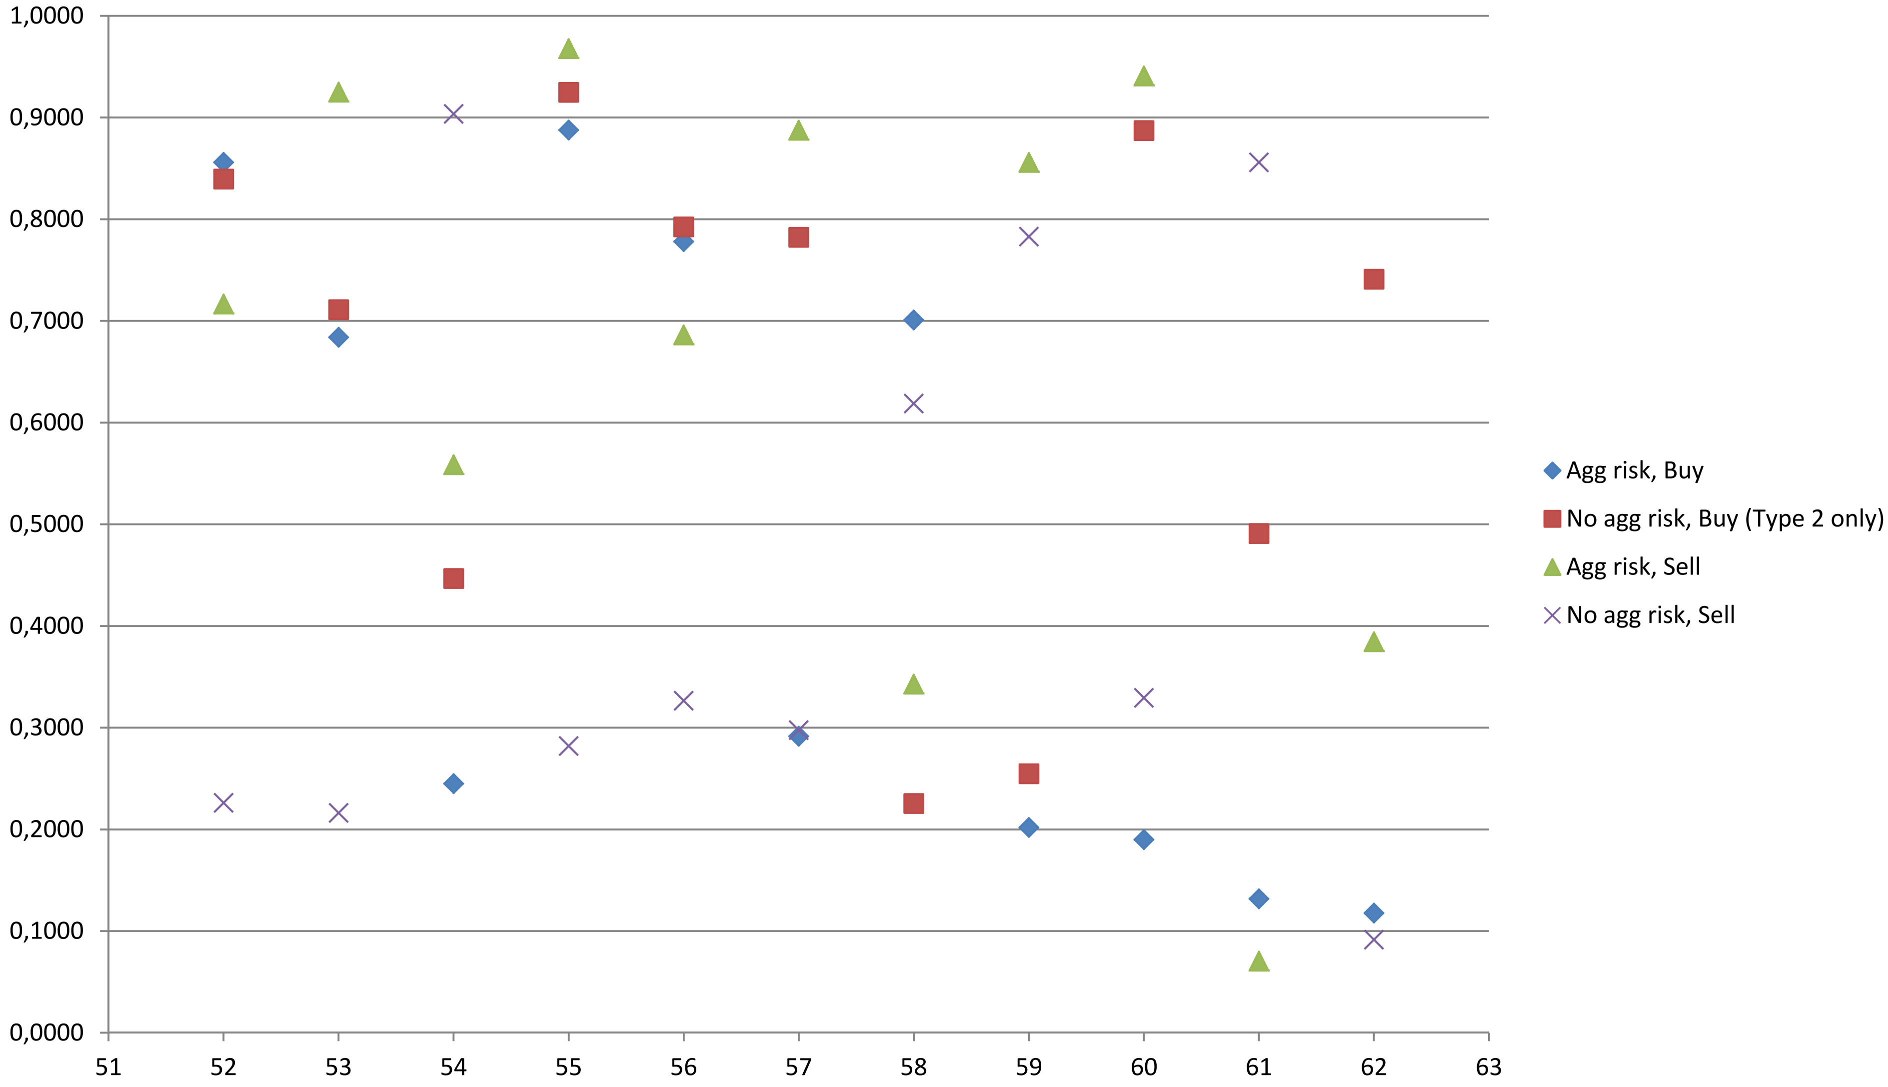

# Figure 6: Percentage of dominated actions averaged across participants

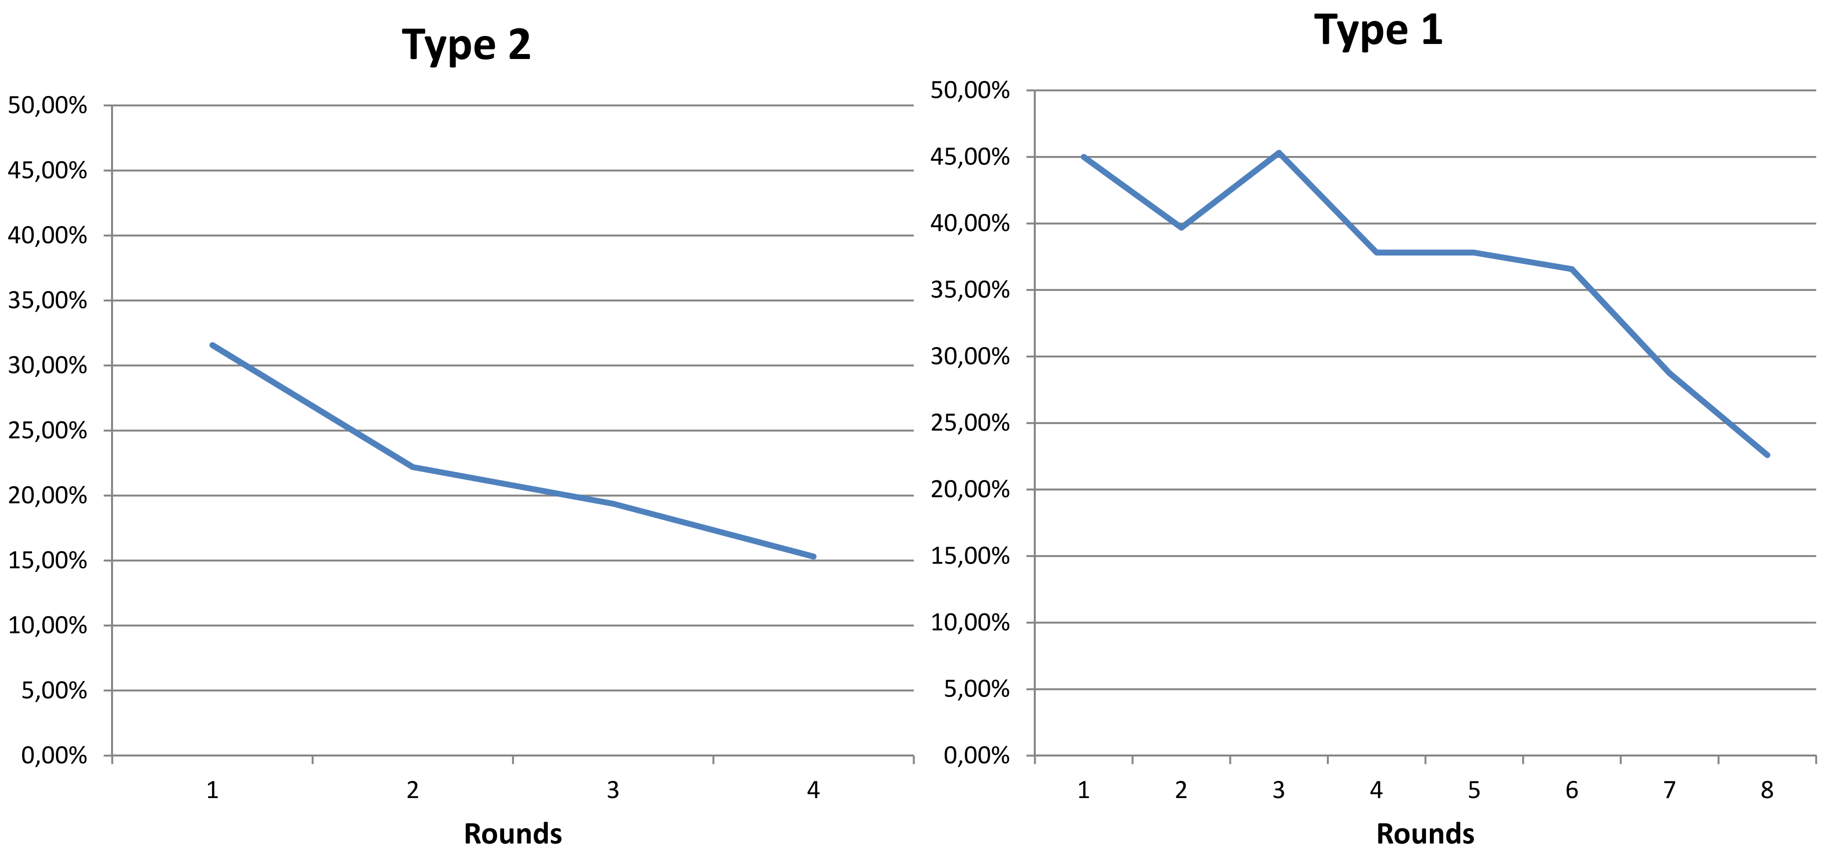

# Figure 7: Excess-supply-minimizing prices in Treatment 1 for the 4 groups

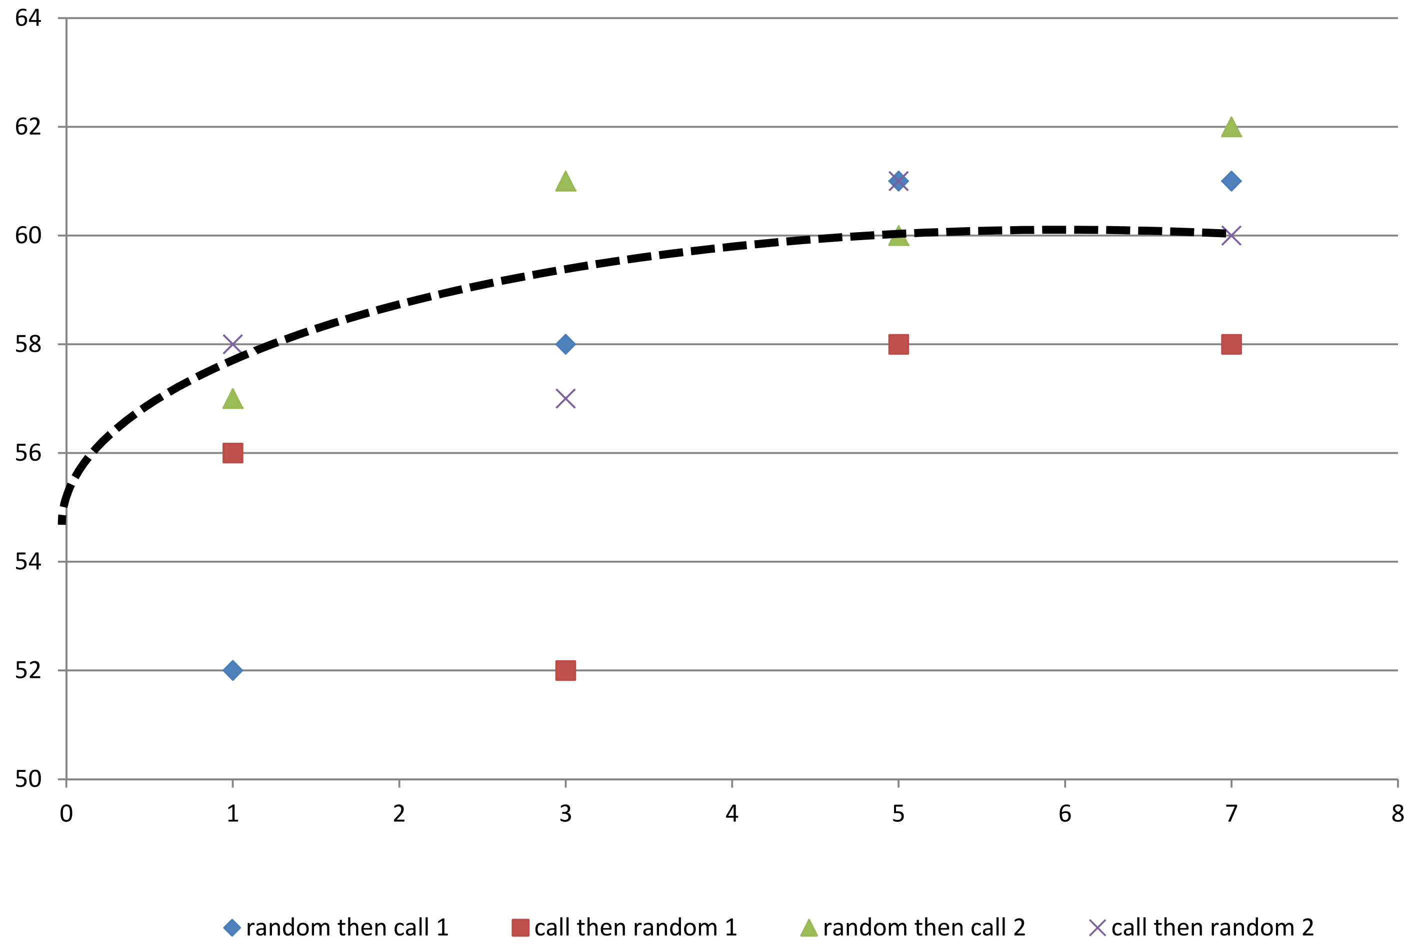

# Figure 8: Excess-supply-minimizing prices in Treatment 2 for the 4 groups

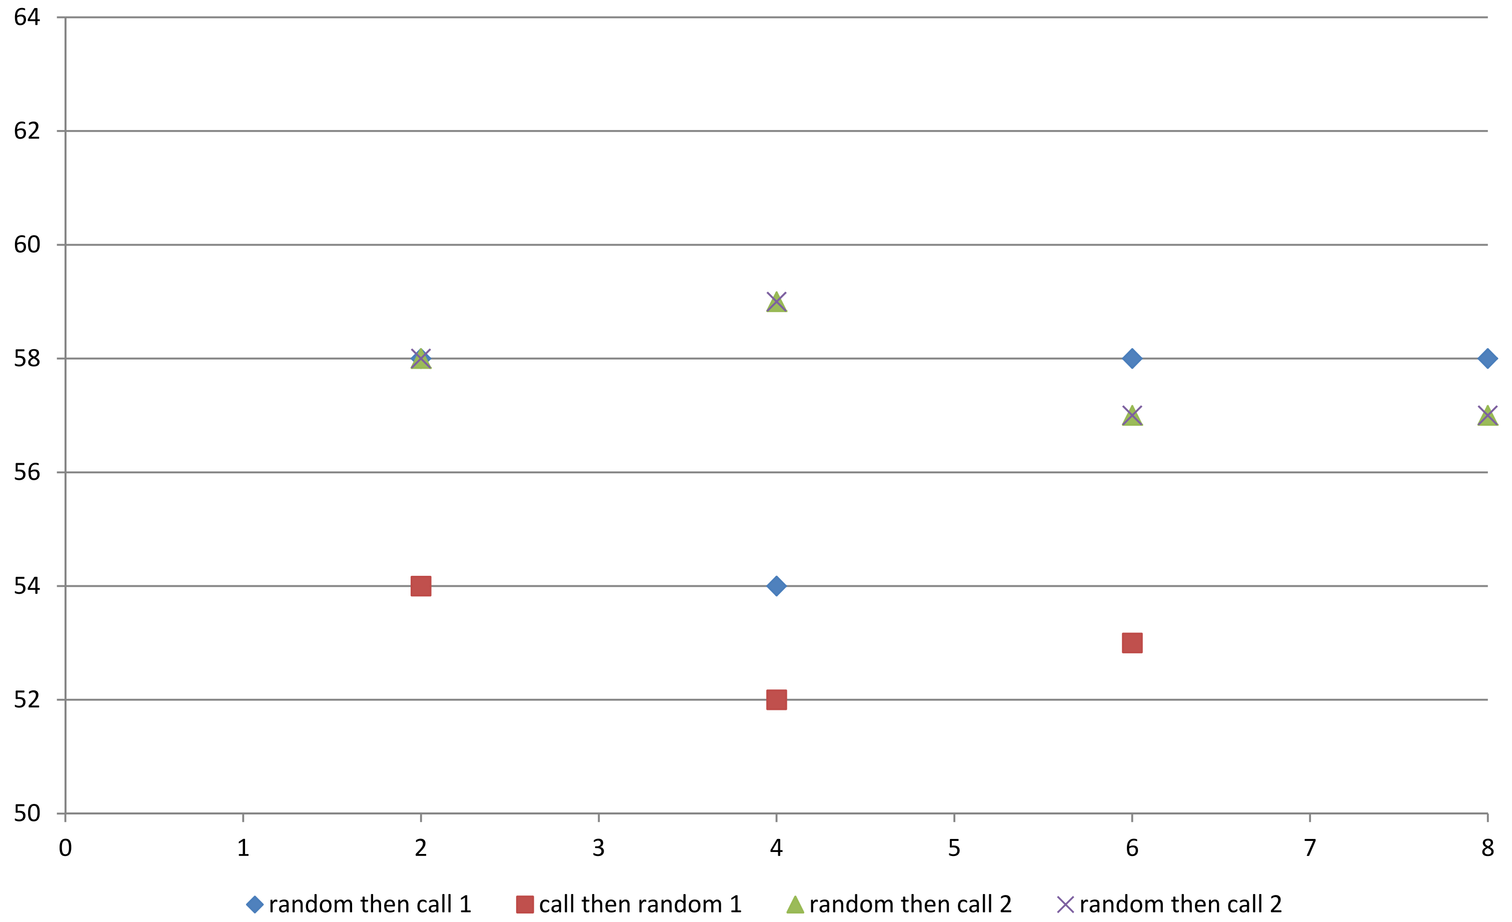

Figure 9: Aggregate supply and demand (over all replications for all groups) in Treatment 1

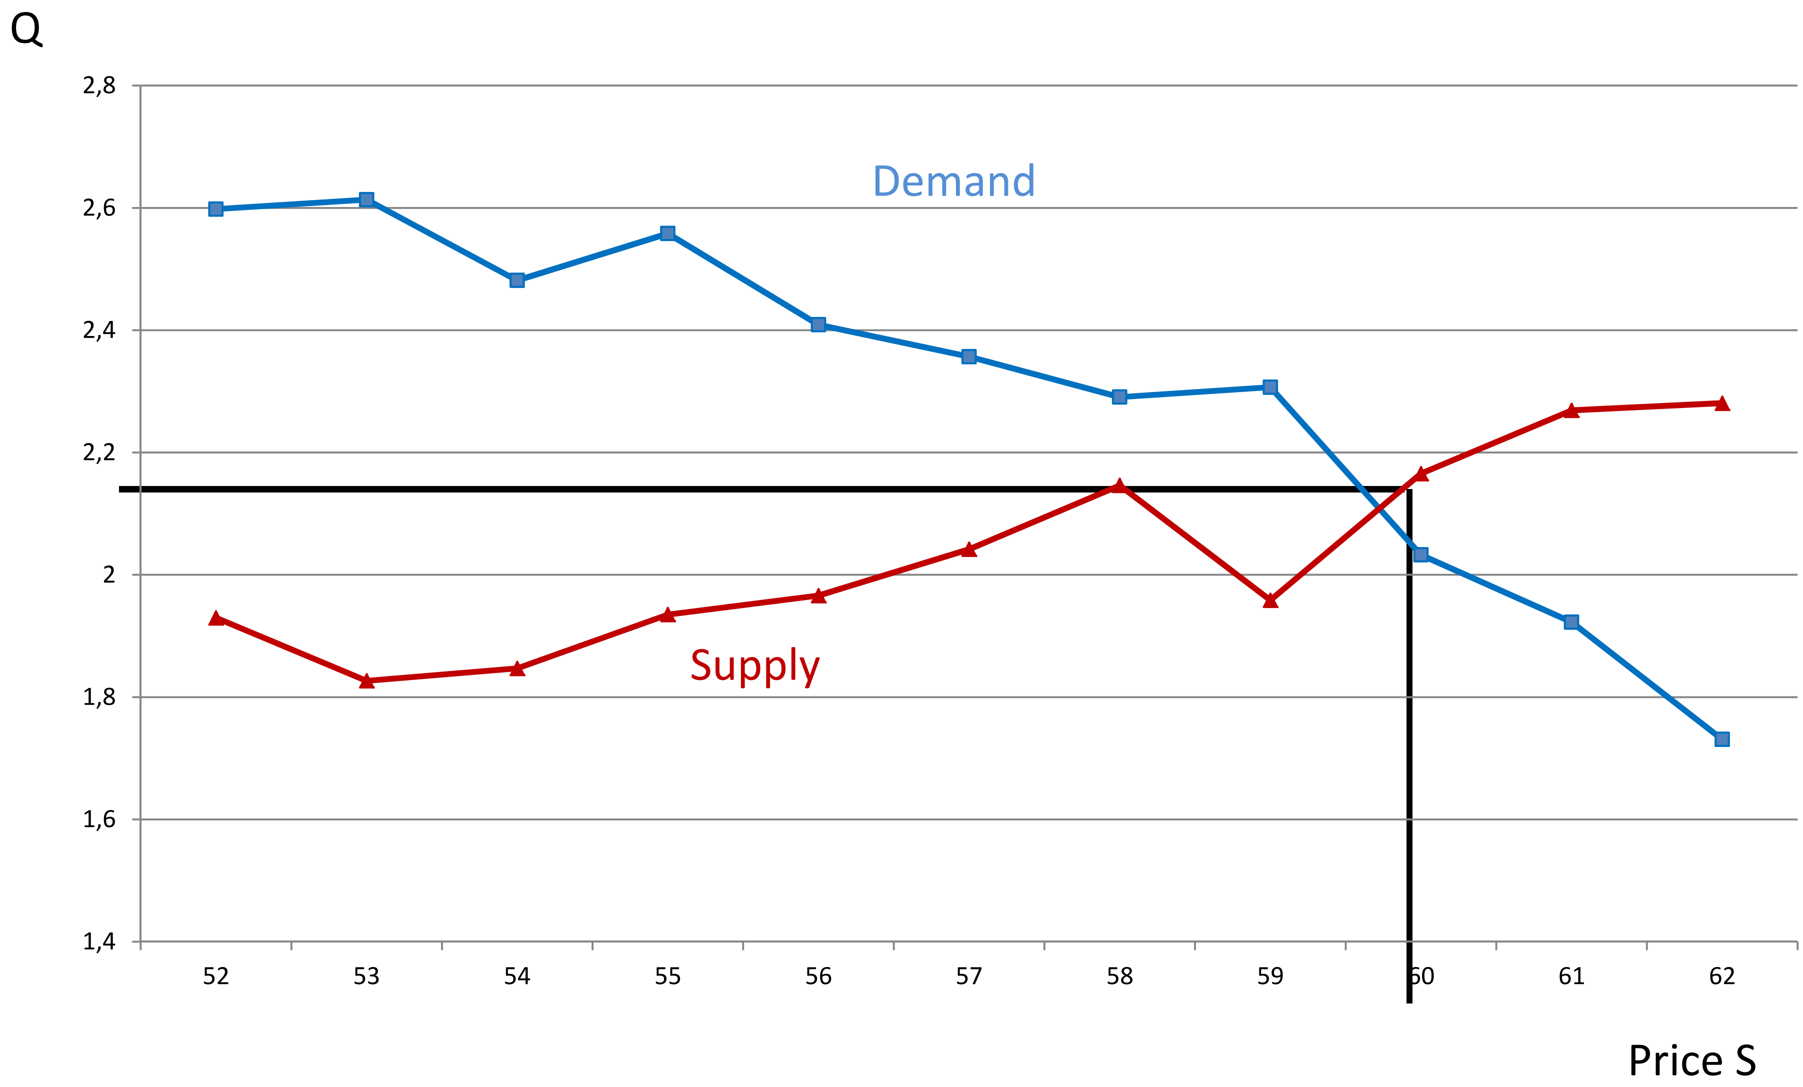

# Figure 10: Aggregate supply and demand (over all replications for all groups) in Treatment 2

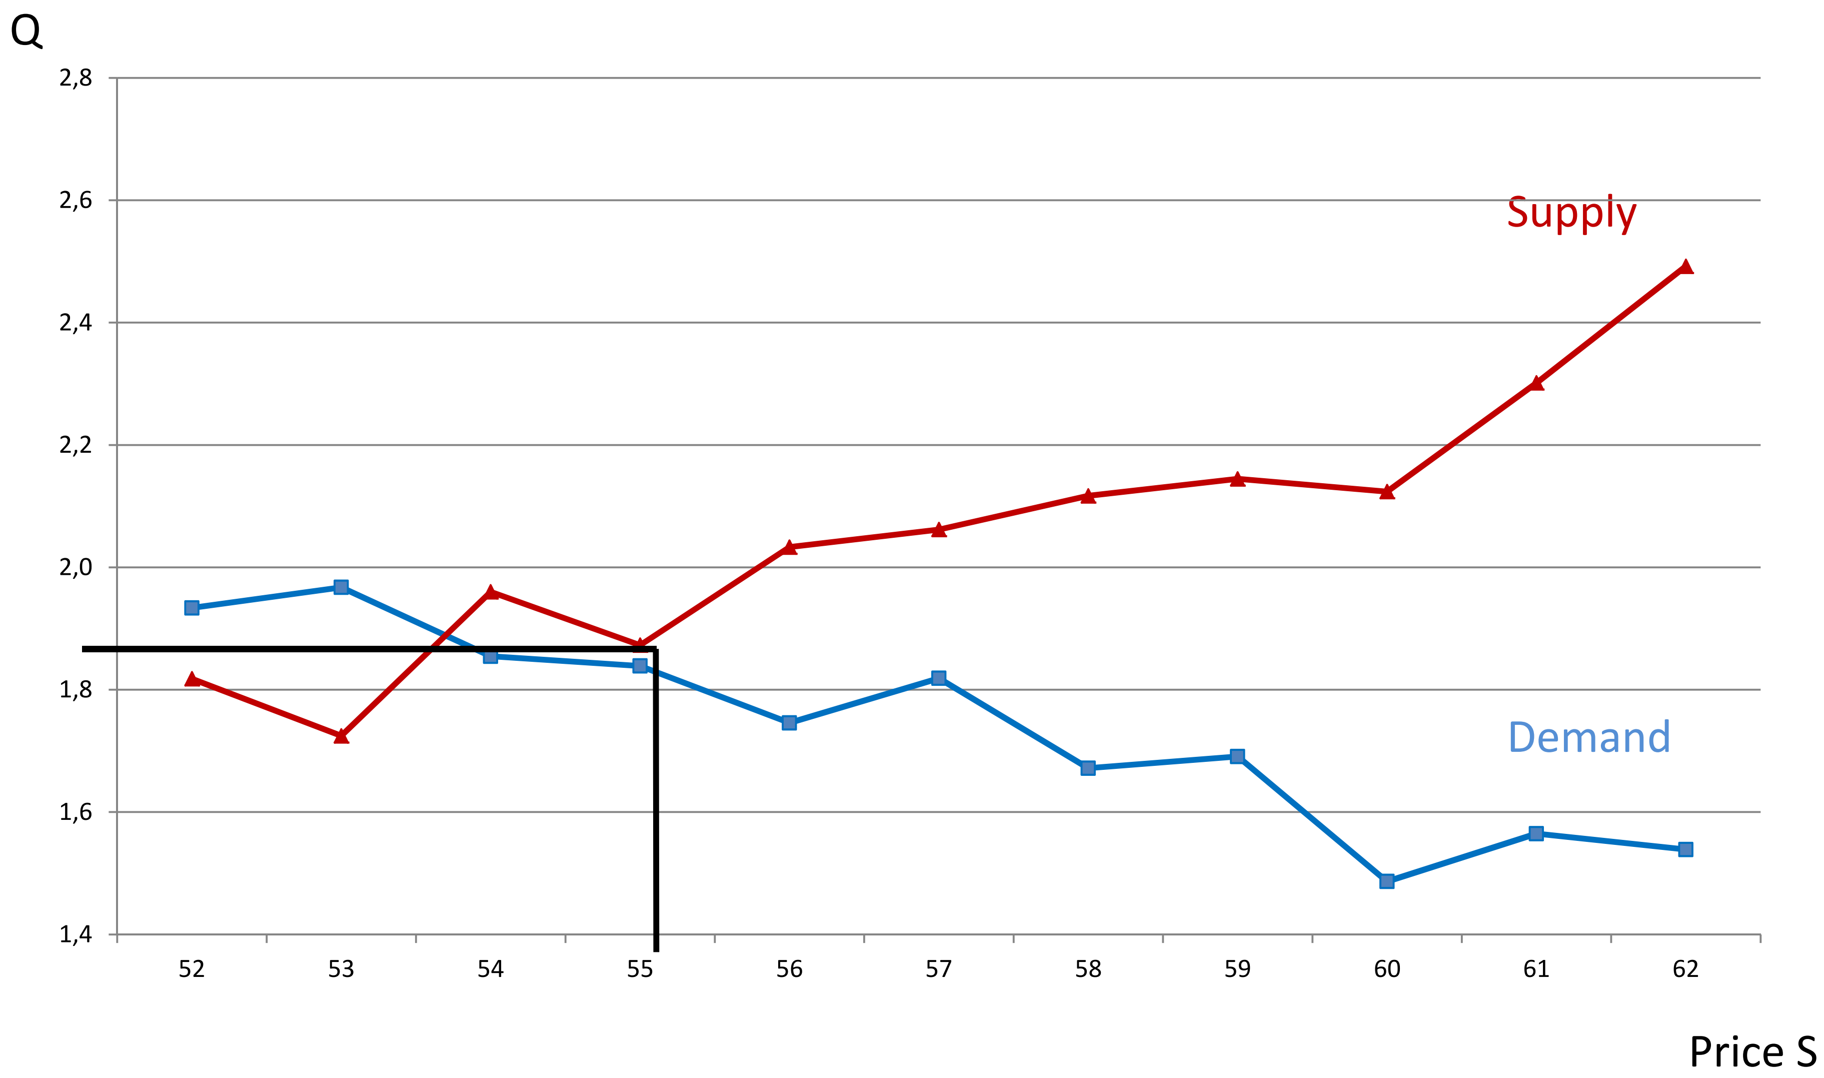

# Figure 11: Distribution of supply and demand in Treatment 1 (no aggregate risk) at overall clearing price *S*=60

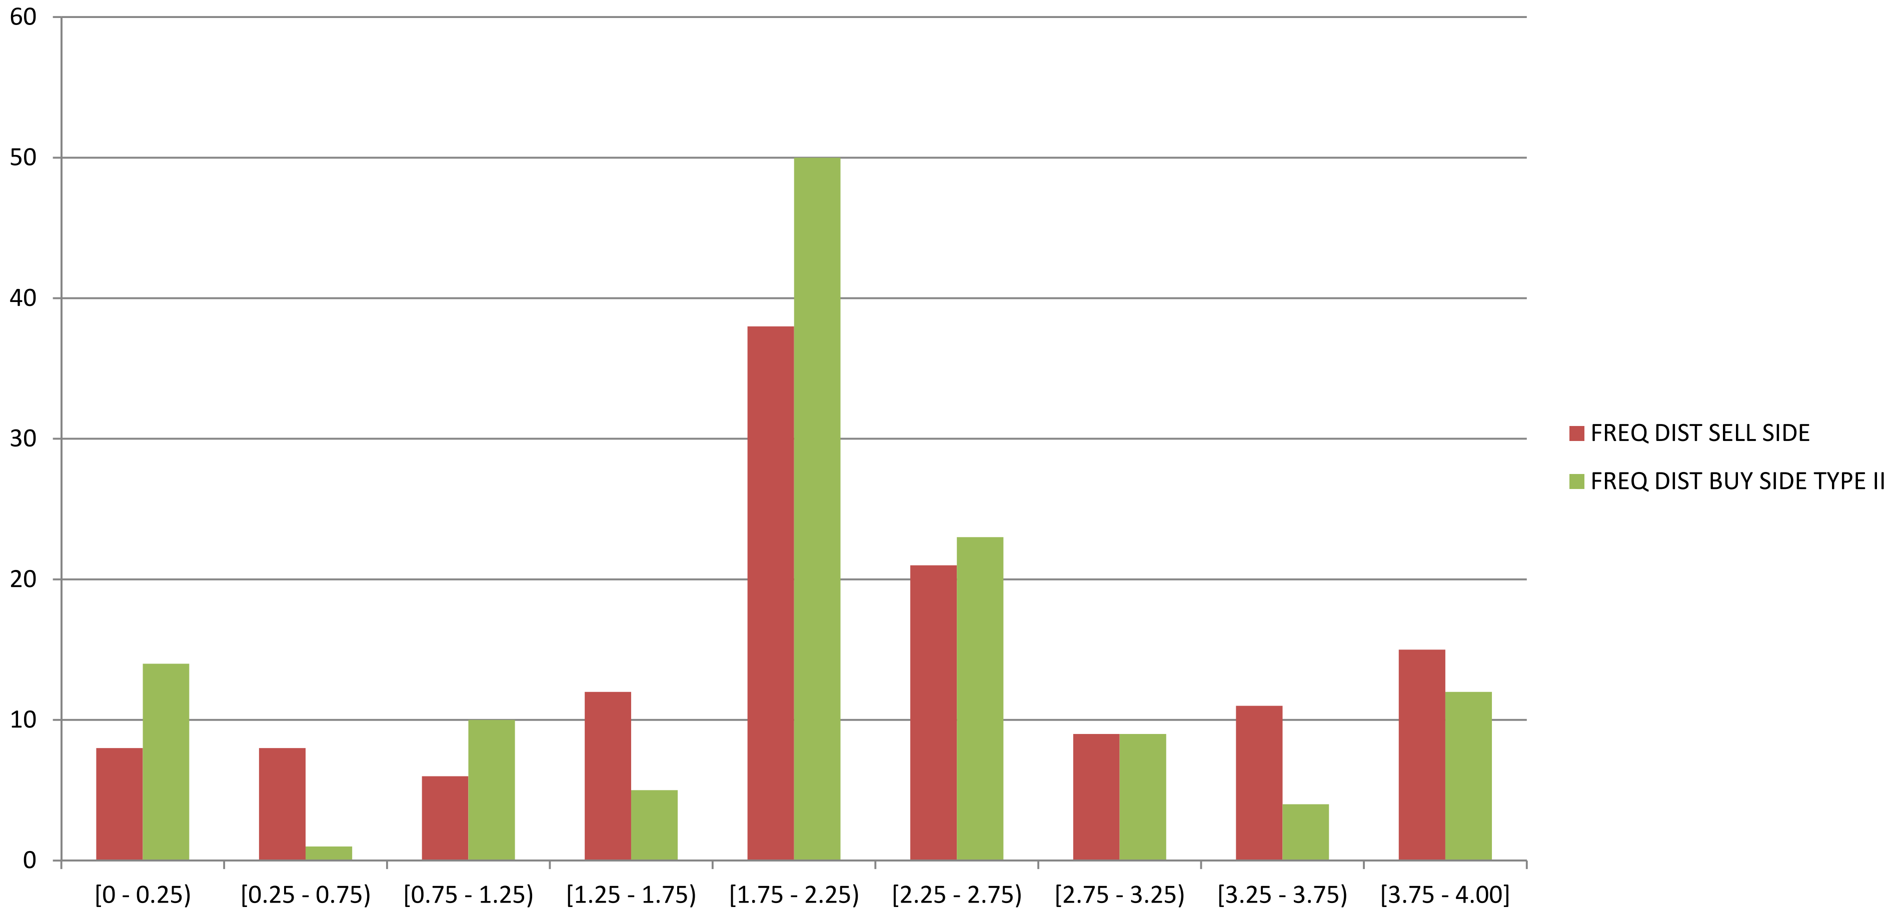

# Figure 12: Distribution of supply and demand in Treatment 2 (aggregate risk) at overall clearing price *S*=55

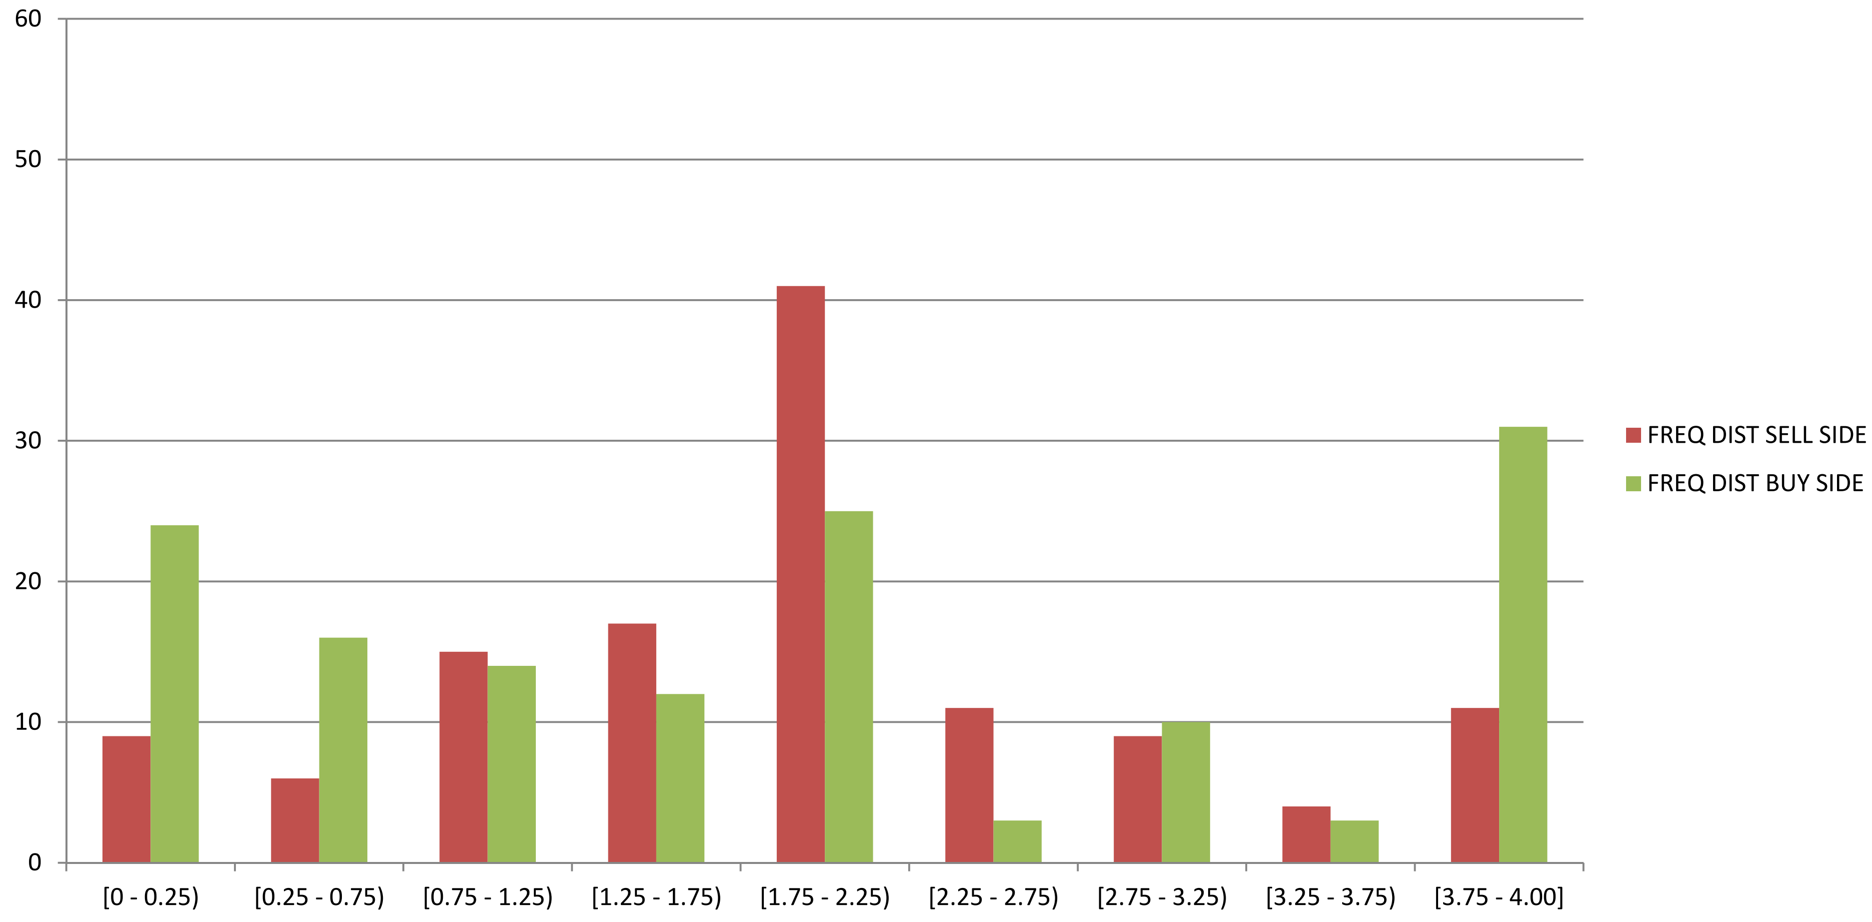

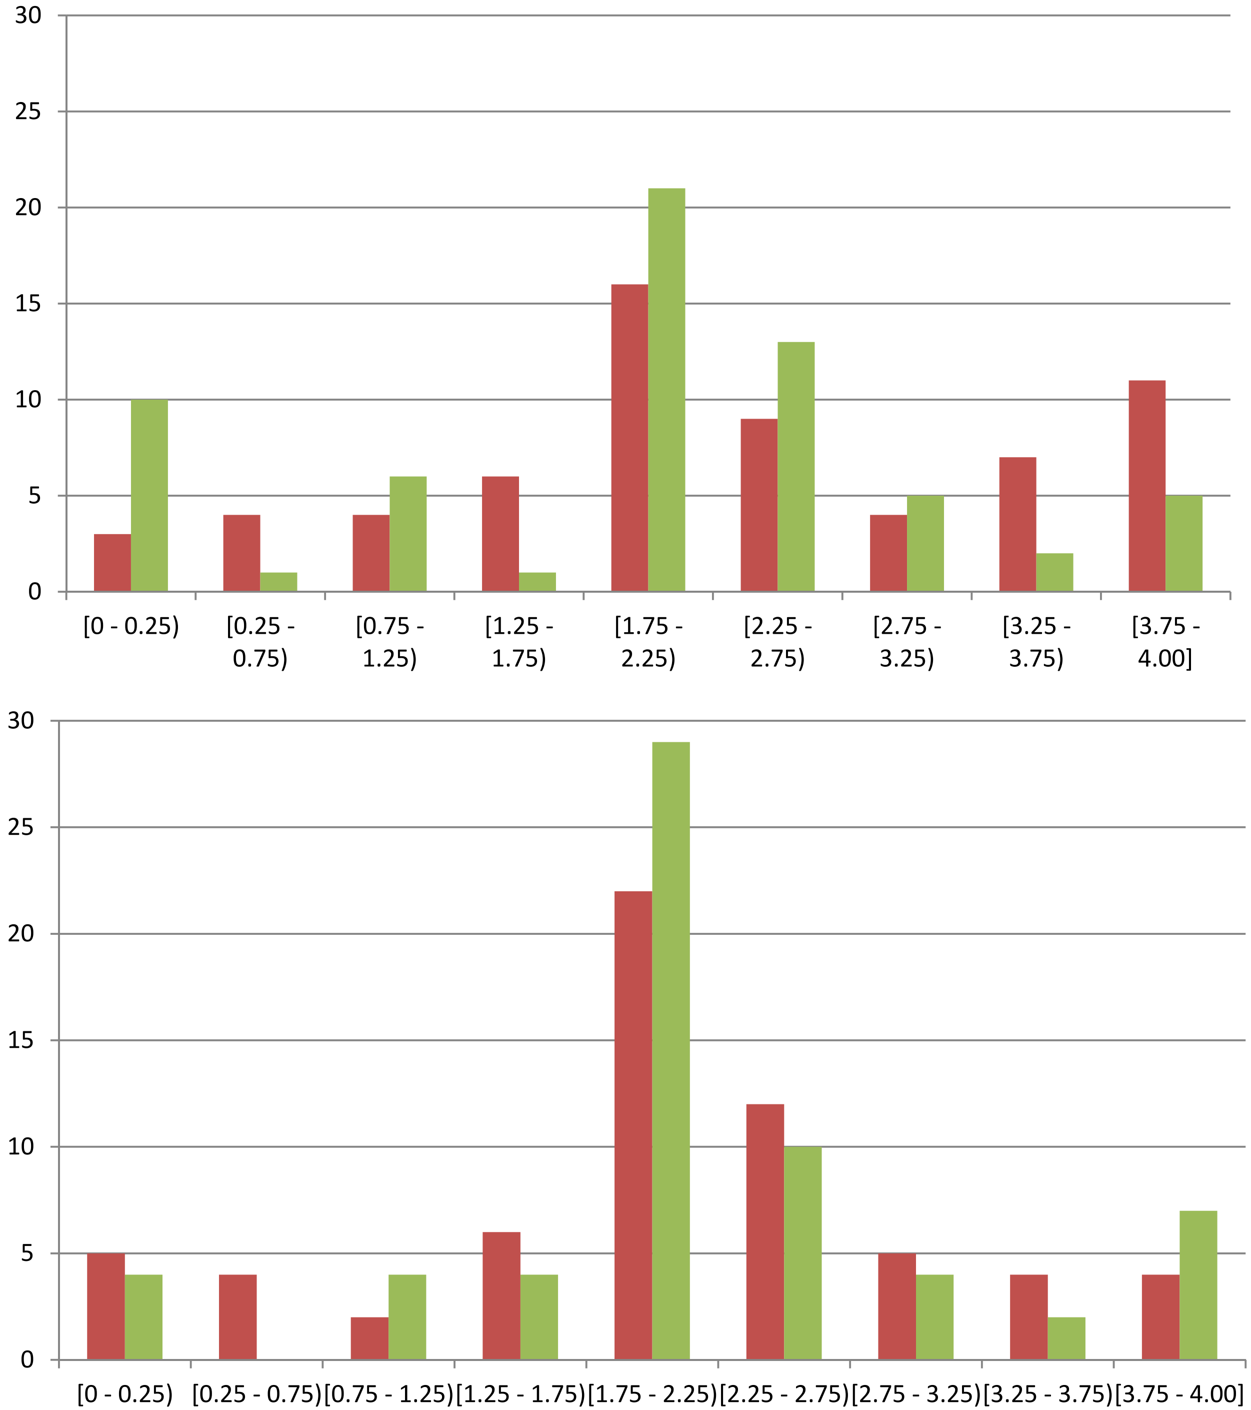


Figure 13: Learning in Treatment 1

Panel A: First 2 replications

Panel B: Last 2 replications

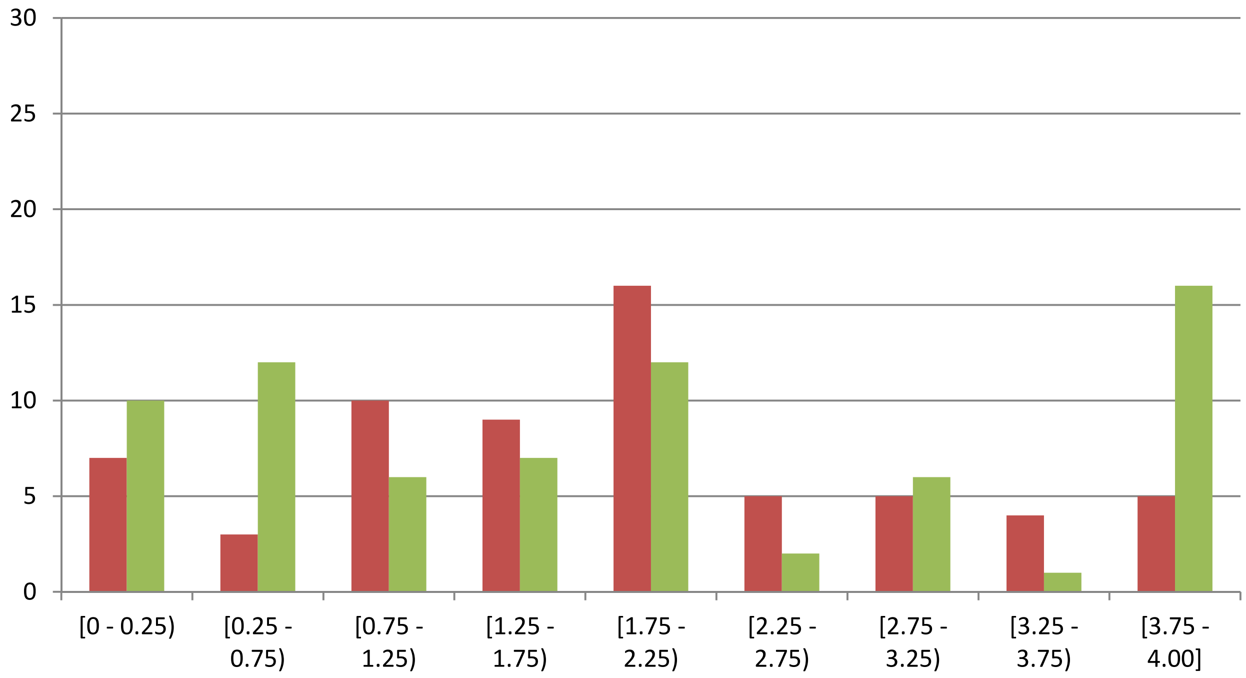


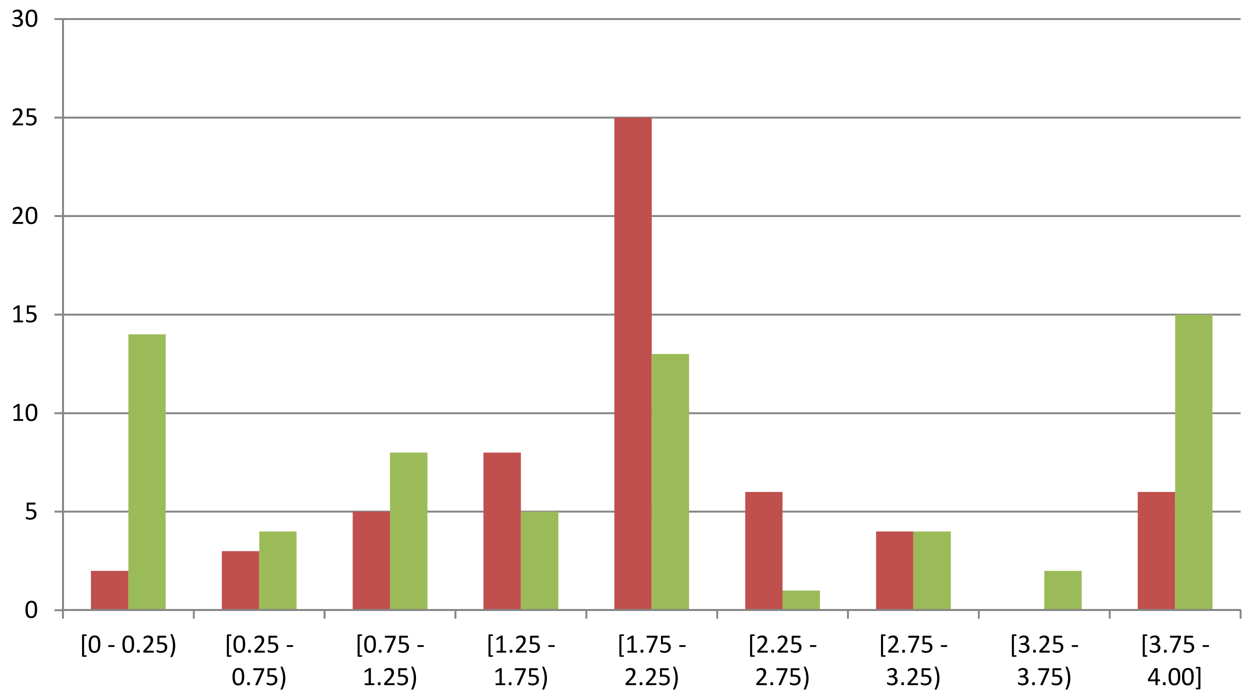


# Figure 14: Learning in Treatment 2

First 2 replications

Last 2 Replications