Pricing in competitive search markets: the roles of price information and fairness perceptions

Nejat Anbarci
Department of Economics
Deakin University
70 Elgar Road, Burwood VIC 3125, Australia
nejat.anbarci@deakin.edu.au

Nick Feltovich*
Department of Economics
Monash University
Clayton VIC 3800, Australia
nicholas.feltovich@monash.edu

Abstract

We use a competitive search (price–posting) framework to experimentally examine how buyer information and fairness perceptions affect market behavior. We observe that moving from 0 to 1 uninformed buyers leads to higher prices in both 2(seller)x2(buyer) and 2x3 markets: the former as predicted under standard preferences, the latter *the opposite of* the theoretical prediction. Perceptions of fair prices – elicited in the experiment – are a powerful driver of behavior. For buyers, fair prices correlate with price–responsiveness, which varies systematically across treatments and impacts sellers' pricing incentives. For sellers, fair prices correlate with under–pricing, which also varies systematically across treatments.

Journal of Economic Literature classifications: D83, D63, E30, C78, C90.

Keywords: directed search; posted prices; frictions; information; price responsiveness; fairness.

*Corresponding author. Financial support from the Australian Research Council (DP140100274), Deakin University and Monash University is gratefully acknowledged. We thank three anonymous referees, participants at the 2014 Singapore Search and Matching Workshop, Kenneth Burdett, Gabriele Camera, Pedro Gomis-Porqueras, Benjamin Lester and Randall Wright for helpful suggestions and comments.

1 Introduction

A fundamental question in search and information economics concerns how the quality of buyers' information about prices affects those prices. As the rise of the internet has led to a proliferation of price–comparison websites for air–fares, hotels, insurance, computers, clothing, and many other goods, the consequences of these websites – and the increased price transparency they bring – for these goods' prices are of obvious relevance. This information is also useful to firms for deciding how transparent to make their pricing schemes, and to policy–makers for deciding whether and how to regulate the reporting of prices. Despite this clear importance, the question has received relatively little attention in the economics literature, perhaps because most would consider the answer to be obvious and long–settled: just as intuition would suggest, early theoretical treatments of search (Salop and Stiglitz 1977; Varian 1980) argued convincingly that making buyers more informed about sellers' prices ought to drive these prices down.[1]

Recently, Lester (2011) has contended that this need not be true; making buyers more informed about prices may actually lead to *higher* prices. He analyzes a model in which capacity–constrained sellers post prices that are observed by *informed* buyers but not *uninformed* buyers, before all buyers choose which seller to visit. Informed buyers must balance the potentially higher surplus gained by visiting a lower–priced seller against the higher chance of being able to buy from a higher–priced seller, leading to an optimal level of price–responsiveness that varies with the numbers of sellers, informed buyers and uninformed buyers. Improving buyers' information by replacing an uninformed buyer with an informed buyer has two countervailing effects. On one hand, losing a price–insensitive buyer and gaining a price–sensitive buyer act to *increase* overall price–responsiveness. On the other hand, the larger number of informed buyers means that competition to buy from the low–priced seller increases; that is, the likelihood of being able to buy when visiting this seller falls. This reduces the incentive to chase low prices, so each informed buyer becomes less price–sensitive. This acts to *decrease* overall price–responsiveness. Often, the latter effect dominates, in which case increasing buyer information *lowers* overall price–responsiveness, giving sellers an incentive to raise their prices – rather than lowering them as intuition might suggest.

This paper is an investigation of how prices are affected by the level of buyer price information, both directly via the competitive pressures described above, and indirectly via its effect on buyers' and sellers' perceptions of what prices are "fair" (of which more below). Empirical testing using observational data from the field is tricky, as natural experiments involving buyer information are rare, and confounding factors are likely to be significant. For example, giving buyers in the field information about prices makes it almost certain that *sellers* will also gain access to this information, making collusive pricing easier to verify and enforce. Thus, when an increase in buyer information is found to lead to an increase in prices (as in Albæk et al. 1997), the result may be attributable to collusion rather than buyers' strategic interaction. To circumvent these problems, we use a lab experiment, allowing us to exogenously vary the variables of interest, and to control precisely the information agents receive. We use four of the simplest non–trivial parameterizations:

[1] See Lester (2011), especially his footnote 2, for examples from empirical studies and the popular press.

2 sellers and either 2 or 3 buyers (2x2 and 2x3 markets), with either 0 or 1 uninformed buyers; we name the four resulting cases "2x2–0", "2x2–1", "2x3–0" and "2x3–1". Importantly, *we vary the number of uninformed buyers within–participant* (i.e., the same participant plays an equal number of rounds with 0 and 1 uninformed buyers), removing confounds due to unobserved individual–specific attributes by using each participant as his/her own control.

Our design allows a strong test of the effects of buyer information. If each buyer and seller is motivated only by his/her money payment (*standard preferences*), moving from 0 to 1 uninformed buyers is predicted to lead to a price *increase* in the 2x2 market but a price *decrease* in the 2x3 market. However, there is ample evidence that individuals are motivated by other considerations such as *fairness perceptions*. As discussed in Section 2, notions of fairness often play a role in decision making, and when markets are thin (so that individuals' choices can substantially affect market outcomes) as in our setting, the effects may be detectable at the market level. We show, using a simple theoretical model of fair prices, that having a preference for fair prices can affect buyer and seller behavior in systematic ways. Perceiving a *lower* price as fair makes sellers choose *lower* prices, while it also makes buyers *more price–responsive*, reinforcing the incentive for sellers to choose lower prices (i.e., fairness views are to some extent self–fulfilling). This means that if perceived fair prices are correlated with price information – that is, if giving buyers more or less price information actually changes what prices buyers and sellers view as fair – then fairness perceptions can amplify, reduce, or even reverse the effects that would follow from competitive pressures on their own.

Our experimental results bear this out. Moving from 0 to 1 uninformed buyers in the 2x2 market does lead to higher prices as predicted by standard preferences. It also increases prices in the 2x3 market, *the opposite* of the prediction. Fairness is the most plausible explanation for our results. Within each treatment, elicited fair prices correlate with behavior in the same ways as were predicted by our model of fair prices. They also vary systematically across treatments, with *higher* prices reported as fair when there is one uninformed buyer than when there are none: both in the 2x2 market where this follows the standard–theory predictions (and thus fairness perceptions *reinforce* the effects of competitive pressures) and in the 2x3 market where it goes in the opposite direction (so that fairness *reverses* the effects of competition). Along the way, we also show that alternative explanations for our results can be ruled out, such as risk aversion, order effects, and incomplete convergence.

2 Relevant literature

One relevant strand of the literature involves experiments based on settings of price posting and directed search (Cason and Noussair 2007, Anbarci and Feltovich 2013, Anbarci et al. 2015, Kloosterman 2015). A recent paper by Helland et al. (2016) is worth special mention because they also examine the effect of buyer information on prices, using an underlying setting very similar to ours. Importantly, they report behavior by both sellers and buyers that conforms much more closely to the standard theoretical predictions than what we find. In particular, their prices *decrease* when moving from 0 to 1 uninformed buyers in the 2x3 market,

where ours *increase*. Their experimental design and procedures differ in many small ways from ours, as we will discuss in Section 4.6; however, we will discover little by way of a satisfying explanation for the different results between their experiment and ours.

A second important strand concerns the effects of fairness perceptions on decisions. Early studies were survey–based (Kahneman et al. 1986, Thaler 1985). A typical result was that price increases are viewed as fair when they are seen as responses to cost increases (e.g., when a retailer increases the price of snow shovels because the wholesale price rises) but not when they are seen as opportunistic (e.g., when a retailer increases the price of snow shovels because a snowstorm leads to high demand). This pattern of perceptions has proved to be quite robust, both in economics (see, e.g., Dickson and Kalapurakal 1994; Pirona and Fernandez 1995; Seligman and Schwartz 1997) and in other business fields (Campbell 1999; Xia et al. 2004 and Campbell et al. 2015 in marketing; Tian and Zhou 2015 in accounting).

The first incentivized study of fairness in markets was that of Kachelmeier et al. (1991a), though other strands of the literature focus on fairness in other settings such as bargaining (see, for example, the discussion of ultimatum–game results in Roth's 1995 survey of bargaining experiments), and the strands do overlap somewhat (see Korth 2009 for an investigation of fairness in a market setting built upon the ultimatum game). Kachelmeier et al. considered 5–buyer–5–seller induced–value markets with buyers posting bids via networked computers. (Sellers could accept buyers' bids but could not post offers.) After 10 rounds, there was a change to the "tax regime", followed by another 10 rounds. Under the new regime, the equilibrium price was higher than before, as were sellers' marginal costs and their share of total surplus. Kachelmeier et al. varied the information given to buyers at the beginning of the second 10 rounds: buyers received no information in a control treatment; they were informed that sellers' costs had risen in a "marginal cost disclosure" treatment; and they were informed that the share of surplus going to sellers had risen in a "profit disclosure" treatment. Kachelmeier et al. found that prices converged to the new higher equilibrium level fastest under marginal cost disclosure and slowest under profit disclosure, consistent with the earlier survey–based studies. They also found that participants' reported fair prices (elicited after the last round, for both the first and second 10 rounds) were lowest under profit disclosure and highest under marginal cost disclosure. They attribute their results to a "principle of dual entitlements" (proposed by Kahneman et al. 1986 and Thaler 1985), according to which consumers in markets receive not only "acquisition utility" (the benefit from the good under standard preferences) but also "transaction utility" (an additional benefit or cost based on how justifiable the price is); the latter gives fair–price perceptions a role in market behavior.

Subsequent studies found similar results. Franciosi et al. (1995) modified Kachelmeier et al.'s (1991a) setting so that sellers posted offers instead of buyers posting bids, and Kachelmeier et al. (1991b) modified it to allow both seller offers and buyer bids; both found convergence fastest in the marginal–cost treatment and slowest in the profit treatment, just as in the original study. Ruffle (2000) found that convergence to a high equilibrium price in posted–offer markets was slower when there were fewer buyers (so that each had a greater effect on market demand) and when the seller share of total profit was higher (so that the equilibrium price was more unfair to buyers), due to differences in "demand withholding" (buyers buying less than the

profit–maximizing amount as a protest) and in its effectiveness. Tyran and Engelmann (2005) modified the setting further to allow buyers to boycott the market (after a majority vote to boycott, the market did not operate for one round); they found evidence of both boycotts and demand withholding, though these served only to reduce efficiency, with little effect on seller price choices. Bartling et al. (2014) examined markets where the buyer had the market power, and found that sellers were less willing to engage in costly punishment following a low price when the price was chosen in an apparently procedurally fair way (via competitive auction) than otherwise (dictated by the buyer). Still other studies have found that fairness can matter in markets with moral hazard (Fehr et al. 1993), adverse selection (Renner and Tyran 2004), and where seller costs are based on performance on a real–effort task (Cason et al. 2011).

3 The experiment

The setting for our experiment comes from Lester's (2011) extension of Burdett et al.'s (2001) directed–search environment (see Figure 1). There is one homogeneous, indivisible, perishable good. There are

Figure 1: Sequence of decisions in the market

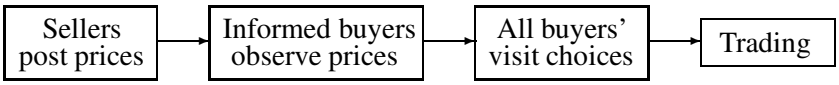

$n \geq 2$ buyers and $m \geq 2$ sellers. Sellers are identical and capacity constrained; the first unit is produced at zero cost but additional production is impossible. Buyers are identical, each with valuation $\bar{U} > 0$ for the first unit and zero for any additional unit. Each seller simultaneously posts a price, which is observed by the $n - u$ informed buyers but not by the u uninformed buyers. Buyers, both informed and uninformed, then simultaneously make their visit choices; each visits exactly one seller. Trade takes place at the seller's posted price; if multiple buyers visit the same seller, one is randomly chosen to buy. Buyers who aren't chosen, and sellers who aren't visited, do not trade.

In the experiment, we set $\bar{U} = 20$; this is with no loss of generality, since a different value of $\bar{U}$ would only re–scale absolute variables such as prices, and would leave relative variables unaffected. The number of sellers m was set to 2, and the number of buyers n to 2 or 3, of which either $u = 0$ or 1 are uninformed. There are thus four experimental cells, which we abbreviate using the form "mxn–u". (E.g., the 2x2–1 cell has 2 sellers, 1 informed buyer and 1 uninformed buyer.) We choose the 2x3 market because the comparison between 2x3–0 and 2x3–1 represents the largest price decrease possible in Lester's framework resulting from changing an informed buyer to an uninformed one, given standard preferences. We choose the 2x2 market because it is the simplest non–trivial directed search market, and because replacing an informed buyer with an uninformed one (moving from 2x2–0 to 2x2–1) has the anticipated effect of a price increase.

3.1 Experimental procedures

The experiment was conducted at Monash University; participants were mainly undergraduates. Participants played 40 rounds with the market (2x2 or 2x3) fixed. There were 20 rounds each with 0 and 1 uninformed buyers; we varied the ordering (0 uninformed then 1 uninformed, or 1 then 0) across sessions.[2] Participants kept the same role (buyer or seller) in all rounds, but were randomly assigned to markets in each round, so as to preserve the one–shot nature of the stage game by having participants interact with different people from round to round. In rounds with uninformed buyers, the uninformed buyer was chosen randomly in each round (so a given buyer would be informed in some rounds and uninformed in others). Some sessions with large numbers of participants were partitioned into two "matching groups" at least twice the size of an individual market, and closed with respect to interaction (i.e., participants in different matching groups were never assigned to the same market), allowing two independent observations from the same session.

The experiment was computerized, and programmed using z–Tree (Fischbacher 2007). All interaction took place anonymously via the computer program; participants were visually isolated and received no identifying information about other participants. In particular, sellers were labeled to buyers within a round as "Seller 1" and "Seller 2", but these ID numbers were randomly re–assigned in each round, and were not made known to the sellers. This anonymity is necessary so that uninformed buyers are truly uninformed, even after many rounds have been played. Also, combined with the random re–assignment into markets each round, this anonymity reduces incentives for repeated–game behavior such as tacit collusion by sellers (by making it impossible to recognize a deviator in future rounds) or dynamic coordination by buyers (e.g., alternating who visits the lower–priced seller). This last point is particularly important because the game played by informed buyers amongst themselves (given sellers' prices) typically has multiple Nash equilibria (see, e.g., Figure 2 in the next section), and it is necessary to remove any coordination devices from the experiment in order to justify use of the symmetric mixed–strategy equilibrium in formulating predictions rather than any of the asymmetric pure–strategy equilibria (which often payoff–dominate the mixed equilibrium).

Written instructions were distributed before the first round, with a small additional set between rounds 20 and 21, and were read aloud in an attempt to make them common knowledge. Sellers began each round by entering their prices, which could be any multiple of AUD 0.05. These prices were displayed to informed buyers, who were also reminded of the number of informed buyers, while uninformed buyers were shown a message stating that they were not being shown the prices. Buyers then made their visit choices, after which participants received end–of–round feedback that included the prices and visit choices in their market, along with an own–profit calculation. Participants received no information about outcomes in other markets.

After the last round, participants were given two additional tasks. First, they completed an Eckel–

[2] Thus the market was varied between–participants, while the number of uninformed buyers was varied within–participants. The use of within–participant variation allows us to control for participant–specific characteristics such as risk aversion when assessing the effect of buyer information. In Section 4.6 we report robustness results based on new sessions with participants playing 50 rounds under a single set of parameters, so that the variation in the number of uninformed buyers is between–participants.

Grossman (2008) lottery–choice task to measure their level of risk aversion. Next, they completed a short questionnaire made up of two questions eliciting participants' views of what a "fair" price would be in rounds 1–20 and what it would be in rounds 21–40 ("In your opinion, what would be a fair price in Part 1 [2] of the experiment?"), followed by several demographic questions.[3] After completing these tasks, participants were paid. Participants received (exactly) the sum of their profits from four randomly–chosen rounds, plus the earnings from the lottery–choice task, plus a show–up fee of AUD 5 for sellers and AUD 15 for buyers.[4] Total earnings averaged about $43 for a session that typically lasted about 60 minutes.

3.2 Equilibrium predictions under standard preferences

For most parameterizations of our model (i.e., values of m, n and u), sellers' and informed buyers' decision problems reflect trade–offs arising from two features of the setting: (1) sellers' capacity constraints, and (2) buyers' simultaneous visit choices. (Uninformed buyers simply visit each seller with equal probability.) Simultaneous visit choices imply that some sellers may be visited by no buyers and thus unable to sell, while others may be visited by multiple buyers, in which case capacity constraints imply that all but one of those buyers is unable to buy. When not all seller prices are the same, informed buyers face a trade–off between earning a large profit conditional on being able to buy (by visiting a low–priced seller) and having a high probability of being able to buy (by visiting a high–priced seller). Usually, the result of this trade–off is that in the game played by informed buyers against each other given the sellers' posted prices (see Figure 2 for an example), the only symmetric equilibrium will be in mixed strategies, with buyers price–responsive but not perfectly so. That is, even higher–price sellers will be visited with positive probability, but lower

Figure 2: The game buyers with standard preferences play in the 2x2–0 cell, given seller prices p_1 and p_2

		Buyer 2	
		Visit Seller 1	Visit Seller 2
Buyer 1	Visit Seller 1	$\frac{20-p_1}{2}, \frac{20-p_1}{2}$	$20-p_1, 20-p_2$
	Visit Seller 2	$20-p_2, 20-p_1$	$\frac{20-p_2}{2}, \frac{20-p_2}{2}$

probability than lower–price sellers.

Given that informed buyers are price–responsive, sellers face a trade–off between earning a large profit

[3] See the appendix for sample instructions and screen–shots, including those for the lottery–choice task and questionnaire. Other experimental materials including the raw data are available from the corresponding author upon request.

[4] At the time of the experiment, the Australian and US dollars were roughly at parity. The differing show–up fees were meant to compensate buyers, who have lower expected earnings. The amounts of the show–up fees, and the fact that they differ between buyers and sellers, were not made known until after the last round, in an attempt to minimize demand effects. Also, participants were not informed of the rounds randomly chosen to determine earnings, nor the realization of the lottery–choice task, until after all decisions had been made, in order to minimize wealth effects. Finally, we acknowledge that our questionnaire, including the elicitation of "fair" prices, was not incentivized, though incentivizing such responses is a non–trivial problem.

conditional on being able to sell (by posting a high price) and having a high probability of being able to sell (by posting a low price). This trade–off usually leads to an interior solution, with sellers posting prices strictly between 0 and buyer valuation in equilibrium. It is worth emphasizing that because the seller earns the posted price if visited by at least one buyer, and zero otherwise, the equilibrium price follows directly from buyers' price–responsiveness – through its effect on the probability of being visited by at least one buyer – so factors affecting price–responsiveness ought also to affect the equilibrium price.

Replacing an uninformed buyer with an informed buyer (i.e., increasing buyer price information) has two effects that pull in opposite directions. The *direct effect* comes from a price–unresponsive uninformed buyer being replaced by a price–responsive informed buyer; this on its own ought to increase overall price responsiveness; that is, it would raise the probability that the low–price seller is visited by at least one buyer. The countervailing *indirect effect* arises from informed buyers' awareness that the number of informed buyers has increased. If individual buyers' price–responsiveness had then stayed unchanged, the new larger number of informed buyers would intensify competition for the low–price seller's item, reducing the expected payoff to visiting that seller relative to the high–price seller. Hence, informed buyers will react by increasing the probability of visiting the high–price seller; that is, they will become less price–responsive.

The overall effect of increasing buyer price information – and in particular, its sign – thus depends on the relative strengths of the direct and indirect effects. Both of the important possibilities are seen in Table 1, which shows the theoretical predictions for each experimental cell under standard preferences.[5] In the 2x2 market, moving from the 2x2–0 to the 2x2–1 cell means that one of the price–responsive informed buyers is replaced by a price–unresponsive uninformed buyer (the direct effect), while the sole remaining informed buyer becomes more price–responsive (the indirect effect). The *direct effect* dominates in this case, so overall price responsiveness decreases; this lower price responsiveness gives sellers an incentive to raise their prices. Thus, moving from the 2x2–0 to the 2x2–1 cell (a decrease in buyer information) is associated with an *increase* in the equilibrium price, as intuition would suggest.

In the 2x3 market, moving from the 2x3–0 to the 2x3–1 cell means – just as in the 2x2 case – that a price–responsive informed buyer is replaced by a price–unresponsive uninformed buyer (the direct effect).

[5]Note that the table shows two sets of point predictions, one based on infinitely–divisible money, and one based on a smallest money unit of 0.05 (as in the experiment). The predictions for the 2x2–0, 2x3–0 and 2x3–1 cells with continuous money are from Lester (2011). In the 2x2–1 cell, the lone informed buyer does not face the trade–off described above, since there are no other informed buyers, and therefore strictly prefers visiting the low–price seller. This makes sellers' profit functions discontinuous: under–cutting the rival by even a cent means being able to sell with certainty (since she will attract the informed buyer), but pricing even a cent higher than the rival means having to hope an uninformed buyer visits. The discontinuous profit function leads to mixed–strategy equilibrium play by sellers, as described in detail in Appendix A.1. The equilibria with discrete money were calculated using the Gambit game theory software (McKelvey et al. 2014), and are available from the corresponding author; these also involve mixing by sellers in 2x2–1 and pure strategies otherwise. Notice that under discrete money, multiple equilibria are shown for the other three cells, but this multiplicity is trivial in that prices differ from the continuous–money equilibrium prices by no more than the smallest money unit. As usual in experiments, we are less interested in the point predictions themselves than in the comparative statics they imply; these are robust to whether we use the continuous–money or discrete–money equilibria as our source of predictions.

Table 1: Theoretical predictions under standard preferences

Cell	Market	Number of uninformed buyers	Expected posted price: continuous money	Expected posted price: discrete money
2x2–0	2x2	0	10.00	9.95, 10.00, 10.05
2x2–1	2x2	1	13.86	13.86
2x3–0	2x3	0	14.55	14.50, 14.55, 14.60
2x3–1	2x3	1	13.33	13.30, 13.35

Then, the two remaining informed buyers both become more price–responsive (the indirect effect). Here, unlike the 2x2 case, the *indirect effect dominates*, so overall price responsiveness increases, giving sellers an incentive to lower their prices. Thus, this decrease in buyer information is associated with a *decrease* in the equilibrium price – the opposite of what intuition would suggest and the opposite of the corresponding comparative static for the 2x2 market.

3.3 Fairness preferences and their effects

The predictions in the previous section are based on both buyers and sellers having standard preferences (i.e., being motivated by expected money payment only). However, the literature discussed in Section 2 suggests a role for fairness considerations. We asked our experimental participants what they thought of as a fair price in each of the blocks of 20 rounds (i.e., with 0 uninformed buyers and with 1 uninformed buyer), but we did so after all rounds had been played. This procedure, which follows Kachelmeier et al. (1991a), has the advantage of avoiding demand effects in market decisions, but the corresponding disadvantage is that reported fair prices may have been affected by participants' experience in the market setting.[6]

There are many potential mechanisms through which fairness concerns might affect behavior. We make an assumption that, while not completely innocuous, seems plausible. Namely, a fair price is a reference point, and participants dislike prices the farther they are from that reference point, while they continue to value their monetary payoff as usual. This implies a fairly minor change to standard preferences; if p is the actual price and p^* is the individual's perceived fair price, then utility is given by

$$U(p, p^*) = \pi(p) - c(p - p^*),$$

where $\pi(p)$ is the monetary payoff and $c(\cdot)$ characterizes distaste for deviations from the fair price.[7] One

[6]As we will see in Section 4.5, the associations we find between participants' reported fair prices and their behavior – for both buyers and sellers – are fairly robust to whether we use data from all rounds or only early rounds. Since early–round decisions would have been made without a good understanding of either the equilibrium or session averages, this robustness suggests that reported fair prices were not overly dependent on experienced prices.

[7]Note that this specification for utility implements the "principle of dual entitlements" (Kahneman et al. 1986a, Thaler 1985); the first term reflects the buyer's "acquisition utility" and the second term reflects his "transaction utility".

special case of c is a quadratic $\left(c(p-p^*) = \frac{\alpha}{2}(p-p^*)^2\right)$, but our results hold for more general forms of c; other then a few technical conditions, we need only for c to be convex, and zero when $p = p^*$.

When $c = 0$, the utility function reduces to standard preferences. Otherwise, fairness preferences make the utility function steeper when the price is less favorable than the fair price (below p^* for sellers, above it for buyers). For prices more favorable than the fair price, fairness preferences flatten the utility function near p^* and reverse the sign of its slope for sufficiently distant prices from p^*. (See Figure 3 for an illustration.) Note that the disutility of prices away from p^* is incurred irrespective of whether trade actually takes place; it comes from posting a particular price for sellers, or from visiting a seller with a particular posted price for buyers.

Figure 3: Sample utility functions under various strengths of fairness preference

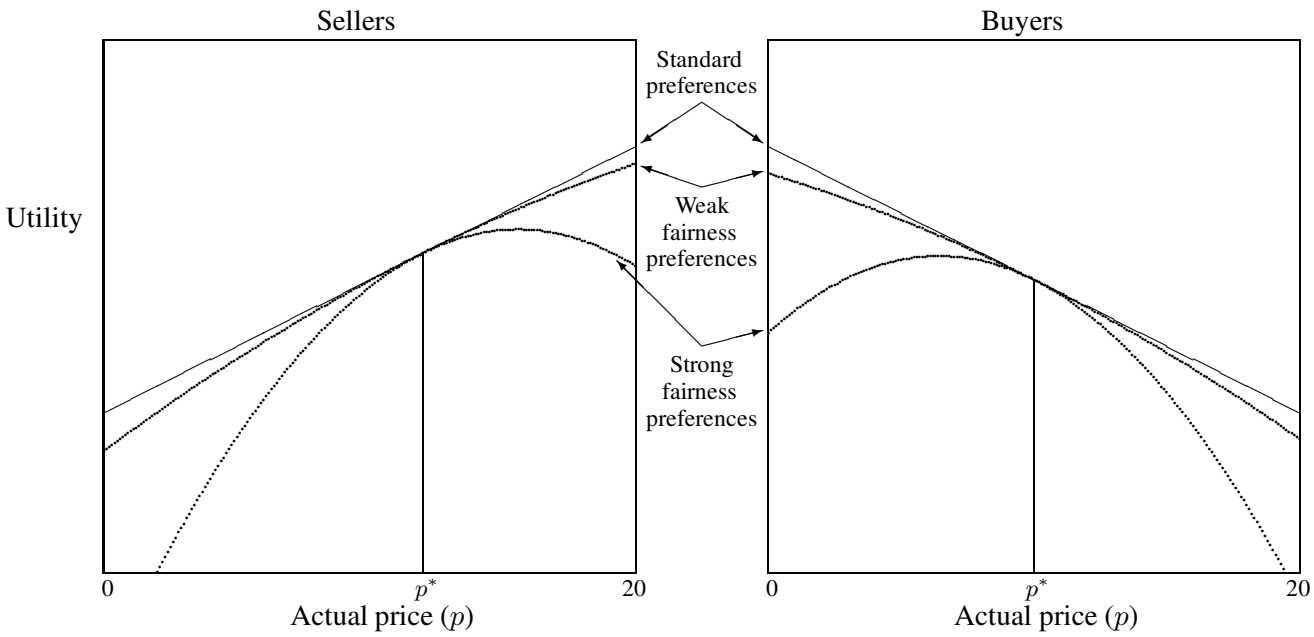

We do not make any assumption about p^* itself: what price(s) individuals actually consider to be fair. In the experiment, seller cost is 0 and buyer value is 20, so participants might view a price of 10 (equalizing the surplus between buyer and seller) as fair. Alternatively, buyers might report lower fair prices than sellers do; since we randomly assign participants to these roles, such a difference would be evidence of a self–serving bias in fairness perception (Babcock et al. 1997). We will take an agnostic approach to the questions of what prices ought to be fair, and whether and how the fair price might differ across treatments. In particular, our analysis of the experimental data will take at face value the fair–price reports we elicit from participants.

Even under the fairly weak conditions we place on c, there is enough structure to show the effect of fair prices on behavior. For informed buyers, lower fair prices lead to increased price responsiveness.

Proposition 1 *Suppose buyers' utility is given by $U(p, p^*) = 20 - p - c(p - p^*)$, where $c(\cdot)$ is twice differentiable, strictly convex, and with $c(0) = 0$.* **(i)** *Suppose all buyers have the same fair price p^*; then*

as p^* decreases, informed buyers become more price–responsive in equilibrium. **(ii)** *Suppose buyers' fair prices are heterogeneous; then informed buyers with lower p^* will be more price–responsive in equilibrium.*

Note that part (i) deals with differences between buyer populations (e.g., across treatments), while part (ii) deals with heterogeneity within populations; the same will be true for the next proposition. Proofs for both parts of both propositions for the 2x2–0 case are shown in Appendix A.2; those for the other cells of the experiment proceed analogously, and are available from the corresponding author upon request.

For sellers, lower fair prices lead to lower posted prices chosen.

Proposition 2 *Suppose sellers' utility is given by $U(p, p^*) = p - c(p - p^*)$, where $c(\cdot)$ is twice differentiable, strictly convex, and with $c(0) = 0$.* **(i)** *Suppose all sellers have the same fair price p^*; then as p^* decreases, the posted price decreases.* **(ii)** *Suppose sellers' fair prices are heterogeneous; then sellers with lower p^* will post lower prices in equilibrium.*

While Proposition 2 is completely intuitive, it is worth remarking on why Proposition 1 holds. Notice from Figure 3 that fairness preferences make the utility function strictly concave. Suppose p_1 and p_2 are the two sellers' prices, and suppose $p_1 < p_2$. (Price responsiveness is not an issue when $p_1 = p_2$, so this assumption is without loss of generality.) In order to visit both sellers with positive probability, it must be that the expected utility from visiting either is the same; that is,

$$(20 - p_1) \cdot \text{Prob}(buy|visit\ 1) - c(p_1, p^*) = (20 - p_2) \cdot \text{Prob}(buy|visit\ 2) - c(p_2, p^*),$$

where $\text{Prob}(buy|visit\ i)$ is the probability of being able to buy if visiting Seller i. Rearranging terms,

$$c(p_2, p^*) - c(p_1, p^*) = (20 - p_2) \cdot \text{Prob}(buy|visit\ 2) - (20 - p_1) \cdot \text{Prob}(buy|visit\ 1); \tag{1}$$

i.e., the difference in "fairness costs" between sellers must equal the difference in expected money payments.

To see the effect of p^* on price–responsiveness, suppose p^* decreases. Because of the convexity of c, the difference $c(p_2, p^*) - c(p_1, p^*)$ increases as a result, so the right–hand side of (1) also has to increase. Since the prices are fixed, it must be that $\text{Prob}(buy|visit\ 2)$ increases relative to $\text{Prob}(buy|visit\ 1)$, which means that the equilibrium probability of visiting Seller 2 must fall relative to that of visiting Seller 1. Since Seller 2 is the higher–price seller, this implies that p^* decreasing has made informed buyers more price–responsive.

A few further remarks about Proposition 1 are in order. First, the key step in the argument uses the convexity of the function c, but the extent of this convexity is not used; thus it is not necessary to have strong fairness preferences (in particular, the result is consistent with monotonically decreasing utility in p), or to have a symmetric cost function (e.g., buyers could dislike prices above p^* more than prices equally far below p^* without losing this result).

Second, while changing the fair price p^* changes price–responsiveness, informed buyers are still typically playing mixed strategies – visiting each seller with positive probability – just as they would under standard preferences. It is possible for the higher–price seller actually to be visited with *higher* probability than the lower–price seller, but this happens exactly where it would be expected: with strong fairness

preferences, and with even the higher posted price well below p^*. (As an illustration, consider the strong fairness preferences for buyers in Figure 3, and suppose both seller prices were close to zero.) Under weak fairness preferences, buyers will visit lower–priced sellers more often than higher–priced sellers – just as under standard preferences – and the same would be true even under strong fairness preferences, as long as the fair price is not too much higher than the sellers' posted prices.

Third, we can use (1) to compare price–responsiveness under fairness preferences with that under standard preferences. Notice that under standard preferences, $c(p_2, p^*) - c(p_1, p^*) = 0$, while under fairness preferences, for $p_1 < p_2$, $c(p_2, p^*) - c(p_1, p^*)$ is negative if both posted prices are less than p^*, and positive if both are greater than p^*. So by similar logic as in the previous paragraph, a buyer with fairness preferences will be more price–responsive than under standard preferences when the fair price is low, and less price–responsive than under standard preferences when the fair price is high.

Fourth, buyers' fairness preferences will affect sellers' behavior. Lower buyers' fair prices mean that buyers will be more price–responsive, giving sellers incentives to lower their prices, and conversely higher buyers' fair prices will incentivise sellers to raise their prices. Thus while buyers' and sellers' notions of fair prices operate through different channels, their separate effects are similar, and their joint effects are as expected: low buyers' fair prices and low sellers' fair prices would be predicted to reinforce each other, as would high buyer and seller fair prices (leading to low actual prices and high ones respectively), while low buyer fair prices and high seller fair prices (or the reverse) should largely cancel each other out.

Finally, our theoretical results, to some extent, capture the spirit of the empirical results involving fairness in markets that were discussed in Section 2. As Kachelmeier et al. (1991a) found in their data, our model predicts associations between fair–price perceptions and specific aspects of market behavior for both buyers and sellers that we will test in our data. Beyond this, previous experiments have few literal implications for ours, since the effect of fairness on our predicted treatment effects depends on how fair–price perceptions differ between markets with 0 and 1 uninformed buyer, whereas in the earlier experiments, it was typically clear – if not explicitly stated – what the relationship between fair and equilibrium prices was. (E.g., in the several experimental designs that increased the equilibrium price part–way through the session, the fair price was taken to be below the new equilibrium price, except perhaps when buyers were given a cost–based explanation for the change.) However, the broad conclusion of these earlier studies – that fairness perceptions can affect the ability of theoretical predictions under standard preferences to characterize observed outcomes – is exactly the implication of our model.

3.4 Hypotheses

The arguments put forth in Sections 3.2 and 3.3 give rise to several testable predictions. We will describe these predictions by first stating two–sided null hypotheses, and then discussing scenarios that would give rise to each of the corresponding directional alternative hypotheses.

Our experimental manipulation involves the number of uninformed buyers in two markets: 2x2 and 2x3.

We begin with the former.

Hypothesis 1 *In the 2x2 market, prices are* equal *with 1 uninformed buyer than with 0 uninformed buyers.*

Recall from Table 1 that under standard preferences, replacing an informed buyer with an uninformed one (moving from 2x2–0 to 2x2–1) should lead to a *rise* in prices. This qualitative effect – based on competitive pressures – is what we would expect to see as long as fairness considerations either play a minor role (e.g., if perceived fair prices are roughly equal between the two cells, or the intensity of fairness preferences is weak) or if they serve to reinforce the effects of competitive pressures (e.g., if perceived fair prices are *higher* in the 2x2–1 cell than in 2x2–0). However, as shown in Propositions 1 and 2, if fairness preferences are strong and work in the opposite direction (i.e., if perceived fair prices are *lower* in the 2x2–1 cell than in 2x2–0), then the opposite treatment effect would be expected: moving from 2x2–0 to 2x2–1 should instead lead to a *fall* in prices.

We continue with the 2x3 market.

Hypothesis 2 *In the 2x3 market, prices are* equal *with 1 uninformed buyer than with 0 uninformed buyers.*

Table 1 shows that under standard preferences – that is, based on competitive pressures alone – replacing an informed buyer with an uninformed one in the 2x3 market should lead to a *fall* in prices. As in the 2x2 case, this is the direction of effect we would anticipate as long as fairness perceptions play a minor role or if they work in the same direction. However, if they work in the opposite direction (i.e., if perceived fair prices are *higher* in the 2x3–1 cell than in 2x3–0) and are sufficiently strong, then the treatment effect could go the opposite way: moving from 2x3–0 to 2x3–1 should instead lead to a *rise* in prices.

Notice that in both markets, we have one alternative hypothesis based on the effect of buyer information coming mostly (or entirely) from competitive pressures, while the opposite alternative hypothesis relies on fairness preferences not only being sufficiently strong, but working in a particular direction. In what follows, we will sometimes refer to these using terminology such as the alternative to a particular null hypothesis "arising from competitive pressures" or "arising from fairness perceptions". However, it is worth remembering that this is to some extent an abuse of terminology: evidence for either alternative hypothesis – or even failure to reject the null – does not mean that either competition or fairness is exclusively driving the results. At most, it might allow a conclusion about the relative strengths of these two channels.

4 Results

We conducted twelve sessions with a total of 206 participants (see Table 2). As noted above, many sessions were partitioned into more than one matching group, which is our independent unit of observation. In total, there were 12 matching groups for each of our 2x2–0 and 2x2–1 cells, and 10 for each of 2x3–0 and 2x3–1.

Table 2: Session information

Cell in rounds		Session	Total number of		
1–20	21–40	numbers	matching groups	markets	participants
2x2–1	2x2–0	2, 9, 12	6	12	48
2x2–0	2x2–1	4, 7, 10	6	12	48
2x3–0	2x3–1	1, 4, 8	4	10	50
2x3–1	2x3–0	3, 6, 11	6	12	60
Totals:			22	46	206

4.1 Market aggregates

In our analysis, we use two price measures – the average *posted price* (the simple average of seller price choices) and the average *transaction price* (the average of those posted prices at which a unit was traded) – though our results are mostly robust to which measure we use. Table 3 reports cell–level posted and transaction prices over all rounds and over the last 5 rounds.[8] Figure 4 disaggregates by matching group, for

Table 3: Aggregate results – prices

Market	Uninformed	Posted price			Transaction price		
	buyers	Predicted	All rounds	Last 5 rounds	Predicted	All rounds	Last 5 rounds
2x2	0	10.00	9.69	9.85	10.00	9.27	9.52
	1	13.86	12.30	12.81	13.33	11.78	12.46
Significance of difference			$p \approx 0.002$	$p \approx 0.007$		$p \approx 0.002$	$p \approx 0.007$
2x3	0	14.55	11.86	12.98	14.55	11.67	12.79
	1	13.33	13.19	14.19	13.33	13.08	14.02
Significance of difference			$p \approx 0.23$	$p \approx 0.08$		$p \approx 0.16$	$p \approx 0.08$

Note: p–values from two–sided Wilcoxon signed–rank tests on matching–group–level data (details in text).

a sense of data heterogeneity, and Figure 5 disaggregates by round to show how prices evolve over time.

Consistent with the "competition" alternative to Hypothesis 1, Table 3 shows that prices in the 2x2 market are higher when there is an uninformed buyer than when both buyers are informed. Figure 4 shows that while there is substantial heterogeneity across matching groups, overall they are roughly centered at the theoretical prediction under standard preferences (shown as a cross), and all but two in each panel show

[8]Note that in the 2x2–1 cell, the predicted transaction price is lower than the posted price. This is because the mixed strategies played by sellers mean that the two realized prices in a market will usually differ, with the lower price more likely to lead to a trade.

Figure 4: Scatter–plots of posted and transaction prices by matching group

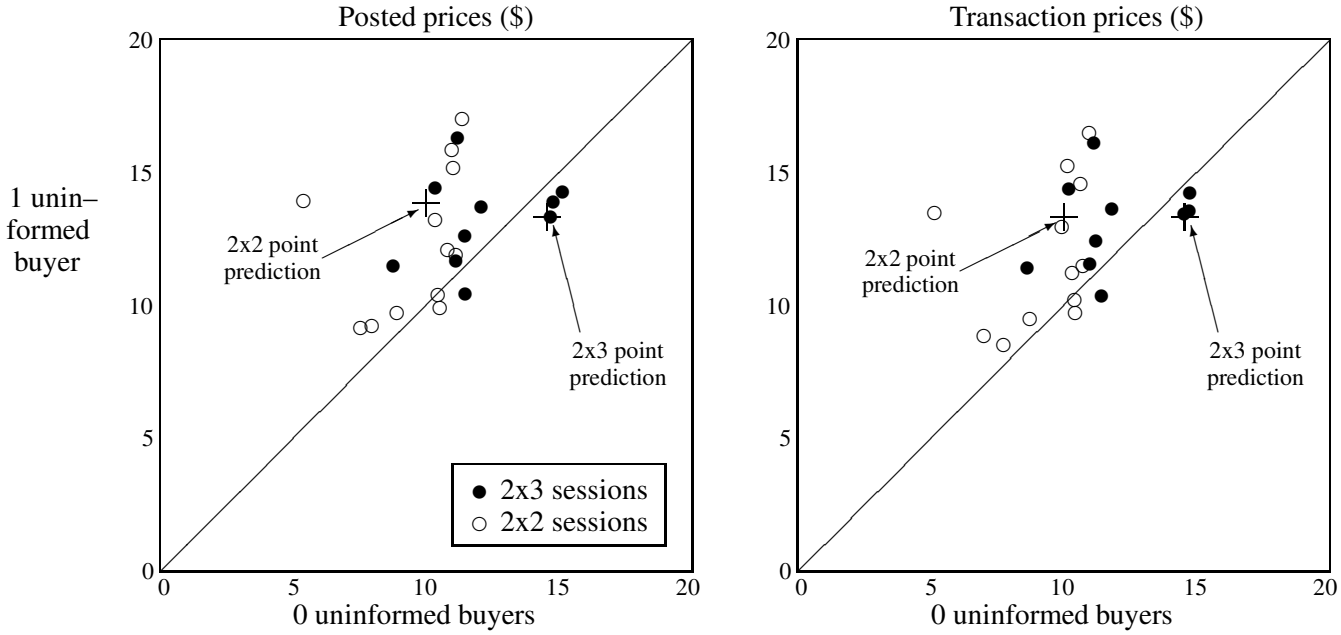

the predicted order relationship between zero and one uninformed buyer. Non–parametric statistical tests on the matching–group–level data verify this apparent treatment effect.[9] Wilcoxon signed–rank tests reject the null hypothesis of no difference in prices between the 2x2–0 and 2x2–1 cells, for both posted prices and transaction prices, and irrespective of whether we use the data from all rounds or from only the last five rounds.[10]

Table 3 shows that prices are also higher in the 2x3–1 cell than in the 2x3–0 cell; this difference is consistent with the "fairness" alternative to Hypothesis 2, and is in the *opposite* direction from what competive pressures on their own would imply. As in the 2x2 market, there is substantial heterogeneity across matching groups in the 2x3 market. In six of the ten groups, the observed difference has the opposite sign as the predicted difference, while in either three (for posted prices) or all (for transaction prices) of the remaining four, the signs are the same but the observed magnitude of the difference is smaller than the predicted magnitude. Non–parametric tests using the matching–group–level data confirm that the differences between 2x3–0 and 2x3–1 cells are significant when the last five rounds' data are used, though not when all rounds are used.

The time series displayed in Figure 5 not only confirm the signs of both treatment effects (2x2–0 versus

[9]See Siegel and Castellan (1988) for descriptions of the non–parametric tests used in this paper, and Feltovich (2005) for critical values for the robust rank–order tests used later in this section.

[10]We use non–parametric tests here, rather than the more common t–test, due to the low number of matching-group–level observations. However, it is worth noting that our significance results are robust to using t–tests instead of Wilcoxon signed–rank tests: while doing so often changes p–values somewhat, there is no case where a significant result in Table 3 becomes insignificant (or the reverse) from using a t–test instead of the Wilcoxon test.

Figure 5: Time series of posted and transaction prices by cell

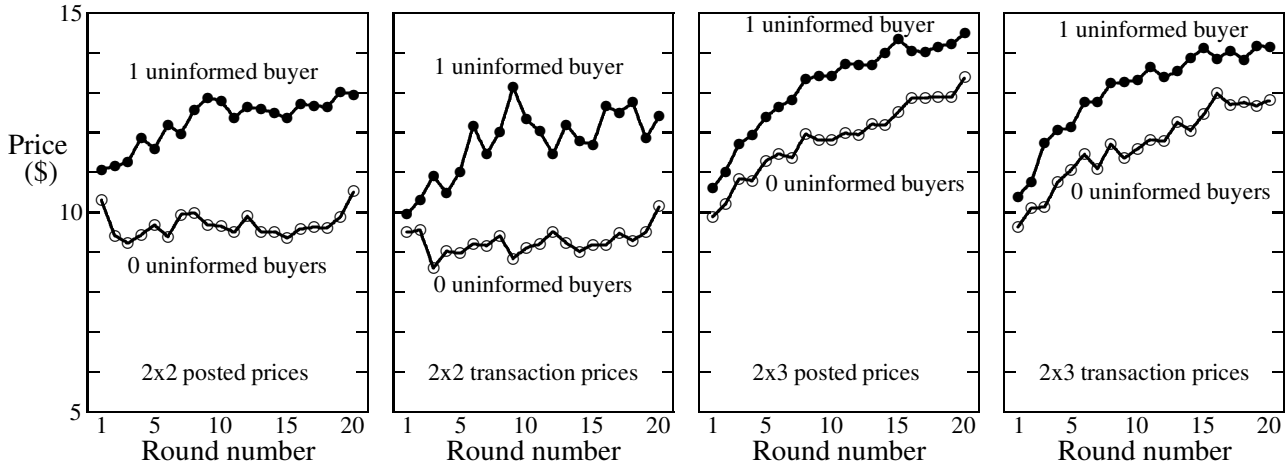

2x2–1 and 2x3–0 versus 2x3–1), but suggest that they are increasing in magnitude over time. So, it is unlikely that the support for the fairness alternative to Hypothesis 2 rather than the competition alternative (or for that matter, the support for the competition alternative to Hypothesis 1) is due to participant inexperience that would have been rectified had more rounds been played.[11]

4.2 Parametric analysis of prices

We next report results from regressions with posted price and transaction price as the dependent variables. Our primary explanatory variables are an indicator for one uninformed buyer and one for the 2x3 market. Additionally, we include the round number (running from 1 to 20, and re–starting at 1 in the second half of the session) and a second–half dummy. Importantly, we include all interactions of these four variables, to allow variables' effects to differ according to other variables' values. (For example, we obviously do not want to force the effect of an uninformed buyer to be the same in the 2x2 and 2x3 markets.) Note in particular that the interaction term between the second–half dummy, the one–uninformed–buyer dummy and the 2x3 dummy allows us to control for order effects in both 2x2 and 2x3 markets.

For both posted and transaction price, we estimate a "restricted" model with only the variables listed above, and an "unrestricted" model with additional variables: the number of sellers in the matching group (as a control for attempts to tacitly collude), the seller's decision time, her self–reported "fair" price, her

[11] Moreover, these results are not due to order effects. If we ignore the data from the second half of the experiment – thus concentrating on the between–participants comparison between those sessions with 0 uninformed buyers in the first half and those with 1 uninformed buyer in the first half – the qualitative results are the same. In the 2x2 market, prices are higher with one uninformed buyer than with none (10.54 and 9.43 respectively for posted prices, and 10.07 and 8.82 for transaction prices), and similarly, in the 2x3 market, prices are higher with one uninformed buyer than with none (12.71 and 10.38 respectively for posted prices, and 12.60 and 10.19 for transaction prices). The observed result for the 2x3 market thus continues to hold, and indeed it is significant despite the sample size being halved (two–sided robust rank–order test, matching–group–level data, $p \approx 0.048$).

choice in the lottery task (as a control for risk aversion), and some demographic variables (sex, age, number of economics classes taken, economics–major indicator, and portion of life lived in Australia). Descriptive statistics for some of these variables (and some that will be used later) are shown in Table 4. We use Stata

Table 4: Descriptive statistics for selected variables used in regressions

Variable	Mean	Std. Dev.	Min.	Max.	Variable	Mean	Std. Dev.	Min.	Max.
Posted price	11.73	3.45	1	20	Transaction price	11.47	3.37	1	20
Seller decision time (sec.)	6.24	7.22	0	76	Number of sellers (per group)	4.29	0.706	4	6
Seller "fair" price	11.73	3.15	5	20	Female	0.481	0.500	0	1
Seller lottery choice	3.67	1.37	1	5	Age	23.03	5.38	18	58
Buyer decision time (sec.)	3.11	4.97	0	76	Economics classes taken	1.13	1.40	0	4
Buyer "fair" price	11.64	3.48	5	20	Economics student	0.097	0.296	0	1
Buyer lottery choice	3.45	1.44	1	5	Portion of life lived in Aus.	0.349	0.391	0	1

(version 12) to estimate panel Tobit models with endpoints 0 and 20 (the bounds on allowable prices).

Table 5 reports estimated average marginal effects and standard errors.[12] The positive and significant marginal effect of the "1 uninformed buyer" variable in both 2x2 and 2x3 markets confirms what was apparent from the treatment aggregates: replacing an informed buyer with an uninformed one is associated with an *increase* in price, both when competitive pressures on their own predict an increase (the 2x2 market) and when they predict a decrease (the 2x3 market). Between the non–parametric tests reported earlier, and the parametric statistical results reported here, the evidence is strong for our main results:

Result 1 *In the 2x2 market, moving from 0 uninformed buyers to 1 is associated with an* increase *in price.*

Result 2 *In the 2x3 market, moving from 0 uninformed buyers to 1 is associated with an* increase *in price.*

Two other variables' effects are worth noting here. The positive marginal effect of the lottery choice is consistent with theory: raising price increases the profit conditional on being visited, but lowers the probability of being visited. So, less risk–averse sellers should choose higher prices. Also, the positive marginal effect of the fair–price variable on actual prices suggests that pricing decisions are influenced by sellers' views about what a fair price is, consistent with Proposition 2.

4.3 Buyer visit choices

In the 2x2–0, 2x3–0 and 2x3–1 cells, the seller prices induce a symmetric subgame played amongst the informed buyers. There are two pure strategies: visit Seller 1 and visit Seller 2. This game always has a

[12] Not shown are the estimates corresponding to the demographic variables, as none of these were significant, though we have left them in Models 2 and 4. Our results are qualitatively robust to (a) using linear instead of Tobit models, (b) estimating separately for the 2x2 and 2x3 markets, (c) estimating based on only the first–half data, (d) combining (b) and (c), or (e) estimating marginal effects in round 20 instead of average marginal effects. These additional results are available from the corresponding author.

Table 5: Tobit results – estimated average marginal effects (unless specified) and standard errors

Dependent variable:	Posted price		Transaction price	
	[1]	[2]	[3]	[4]
1 uninformed buyer...	2.103^{***}	1.634^{***}	1.995^{***}	1.574^{***}
	(0.069)	(0.073)	(0.071)	(0.074)
...in 2x2 market	2.659^{***}	2.007^{***}	2.534^{***}	1.947^{***}
	(0.096)	(0.102)	(0.102)	(0.107)
...in 2x3 market	1.496^{***}	1.229^{***}	1.464^{***}	1.205^{***}
	(0.100)	(0.099)	(0.099)	(0.097)
2x3 market	1.509^{***}	0.950^{**}	1.748^{***}	1.238^{***}
	(0.430)	(0.430)	(0.414)	(0.415)
Round number	0.104^{***}	0.100^{***}	0.114^{***}	0.105^{***}
	(0.006)	(0.006)	(0.006)	(0.006)
Second half	1.975^{***}	1.489^{***}	2.022^{***}	1.515^{***}
	(0.060)	(0.075)	(0.071)	(0.077)
Number of sellers (per matching group)		-0.895^{**}		-0.876^{**}
		(0.393)		(0.379)
Seller decision time (sec.)		-0.011^{*}		-0.028^{***}
		(0.006)		(0.015)
Seller fair price		0.420^{***}		0.398^{***}
		(0.026)		(0.027)
Seller lottery choice		0.419^{***}		0.378^{***}
(higher=less risk averse)		(0.149)		(0.144)
Constant term?	Yes	Yes	Yes	Yes
Demographic variables?	No	Yes	No	Yes
N	3680	3680	2959	2959
$\|ln(L)\|$	8063.83	7931.54	6285.64	6161.87

* (**,***): *Marginal effect significantly different from zero at the 10% (5%, 1%) level.*

unique symmetric Nash equilibrium that specifies the probability of visiting each seller. For example, if p_1 and p_2 are the sellers' prices, each informed buyer should visit Seller 1 with probability

$$q(p_1, p_2) = \frac{20 - 2p_1 + p_2}{40 - p_1 - p_2}. \tag{2}$$

in the 2x2–0 case under standard preferences (see Figure 2 in Section 3.2 for the payoff matrix).[13]

From here, we construct a *reliability diagram*, showing how closely the actual probability of visiting a seller corresponds to this prediction. This is known as the *calibration* of the symmetric equilibrium as a predictor of buyer visit choices. If calibration is high, then over all cases where the equilibrium predicts a given seller is visited with probability 0.4, buyers should actually have chosen to visit that seller about

[13] This expression holds as long as it is defined (i.e., except when $p_1 = p_2 = 20$), and between 0 and 1. When it is larger than 1 (smaller than 0), visiting Seller 1 (Seller 2) is a dominant strategy for each buyer. There typically exist asymmetric pure–strategy equilibria also, but the symmetric one is favored by theorists (Burdett et al. 2001; Lester 2011).

four–tenths of the time; when the predicted probability is 0.7, she should have been visited seven–tenths of the time by buyers; and so on.

To construct the reliability diagram, we first compute the predicted probability of visiting Seller 1 for each buyer and round under standard preferences, and record the buyer's actual visit choice in that round. (Since there are only two sellers, using Seller 2 instead of Seller 1 would yield an equivalent diagram.) Second, for each of thirteen intervals of predicted probability $(\{0\}, (0, 0.1], (0.1, 0.2], (0.2, 0.3], (0.3, 0.4], (0.4, 0.5), \{0.5\}, (0.5, 0.6], (0.6, 0.7], (0.7, 0.8], (0.8, 0.9], (0.9, 1), \{1\})$, the average of the predicted probabilities lying in the interval is calculated, as is the frequency of actual visits to Seller 1 in those occurrences. Finally, at the ordered pair $(average\ predicted\ probability, actual\ frequency)$, a circle or square is plotted with area proportional to the number of occurrences. For example, if there are 100 cases where the predicted probability of visiting Seller 1 was 0.5, and the buyers actually visited Seller 1 in 47 of those cases, a circle is plotted at (0.5, 0.47), with area proportional to 100. The reliability diagrams for all four cells are shown in Figure 6. Also shown are OLS trend lines for each cell, along with the 45–degree line

Figure 6: Buyers' symmetric Nash equilibrium probability of visiting Seller 1 versus actual frequency of visits to Seller 1 *(area of a circle or square is proportional to the number of observations it represents)*

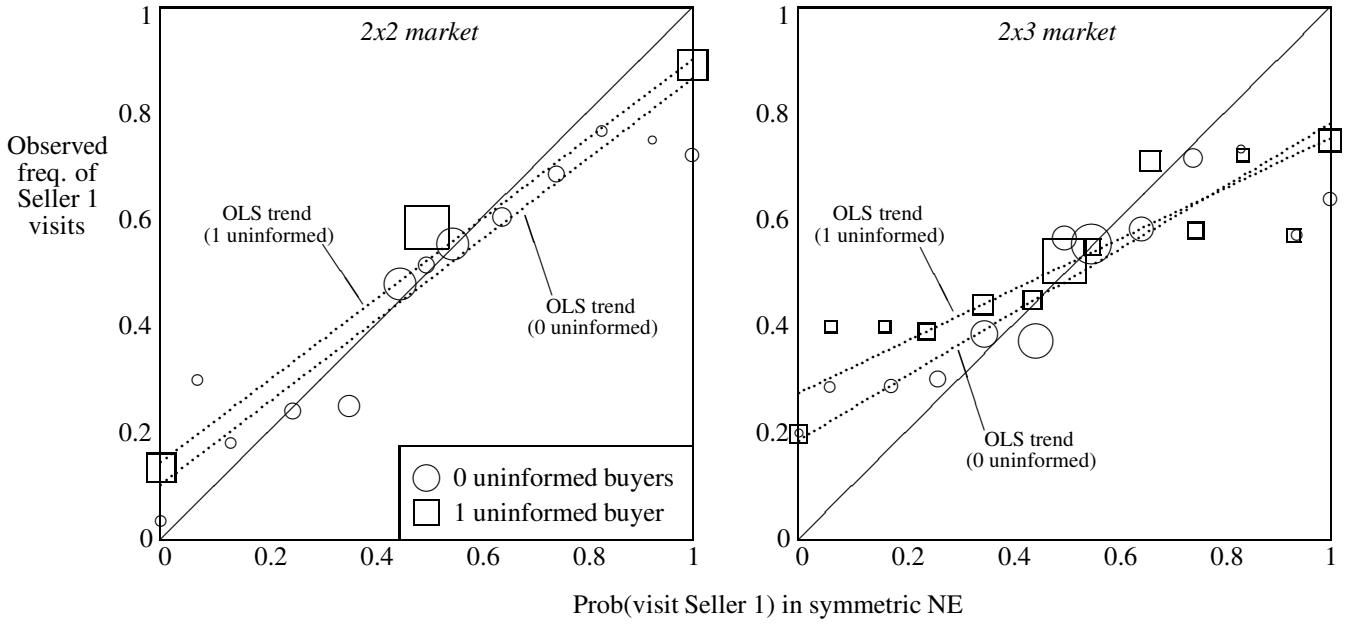

(where predicted and observed probability are equal, and hence calibration is perfect).

Two aspects of buyer visit behavior are apparent. First, buyers are less price–responsive than predicted, as OLS trends are flatter than the 45–degree line. Second, calibration varies systematically across cells. It is difficult to detect any differences between the 2x2–0 and 2x2–1 cells, but buyers are visibly more price–responsive in the 2x3–0 cell than in the 2x3–1 cell. These differences can also be seen in panel probit regressions based on the informed–buyers' decisions, with a Seller–1–visit dummy as the dependent variable. The main explanatory variables are the predicted probability under standard preferences, 2x3–market

and 1–uninformed–buyer dummies, the round number and the second–half dummy, along with all interactions and a constant. A second probit includes these variables and also the buyer's decision time, reported "fair" price, lottery choice, and demographic variables (on their own and interacted with the predicted probability), and a third probit is identical to the second except that it uses only the first 7 rounds of data rather than all 20, as an indication of inexperienced–buyer behavior.[14] The results are shown in Table 6; note that starred values are significantly different from *one* (full price responsiveness) rather than the usual zero.

Table 6: Marginal effect of predicted Seller 1 visit probability on observed Seller 1 visit (informed buyers only), at particular variables' values (panel probits)

	[5]	[6]	[7]
	(All rounds)		(Rounds 1–7)
Overall	0.610^{***}	0.653^{***}	0.653^{***}
	(0.037)	(0.037)	(0.037)
2x2–0 cell	0.734^{***}	0.696^{***}	0.617^{***}
	(0.063)	(0.071)	(0.117)
2x2–1 cell	0.853^{***}	0.851^{***}	0.768^{***}
	(0.042)	(0.044)	(0.178)
Significance of difference:	$p \approx 0.12$	$p \approx 0.06$	$p \approx 0.48$
2x3–0 cell	0.624^{***}	0.681^{***}	0.726
	(0.062)	(0.061)	(0.104)
2x3–1 cell	0.479^{***}	0.552^{***}	0.379^{***}
	(0.051)	(0.049)	(0.093)
Significance of difference:	$p \approx 0.07$	$p \approx 0.09$	$p \approx 0.01$
Lottery choice = 1		0.345^{***}	0.314^{***}
		(0.067)	(0.109)
Lottery choice = 5		0.798^{***}	0.750^{***}
		(0.045)	(0.074)
Significance of difference:		$p < 0.001$	$p \approx 0.001$
Fair price = 8		0.745^{***}	0.712^{***}
		(0.045)	(0.072)
Fair price = 13		0.599^{***}	0.553^{***}
		(0.039)	(0.065)
Significance of difference:		$p < 0.001$	$p \approx 0.01$

* (**,***): *Marginal effect significantly different from one at the 10% (5%, 1%) level.*

The table confirms the under–responsiveness that was seen in Figure 6. While the marginal effect of the predicted probability is positive (indicating that informed buyers do respond to price), it is significantly less than one – overall and in each individual cell. Even more importantly, price responsiveness clearly differs

[14] Ideally, we would use fewer rounds than this, since even seven rounds is sufficient for learning to take place. However, Stata was unable to estimate the regression model when fewer rounds were used, due to lack of concavity of the likelihood function.

across cells. Over all rounds (Models 5 and 6), chi–square tests easily reject the null of equal responsiveness across all four treatments ($\chi^2_3 = 35.52$ for Model 5, $\chi^2_3 = 22.24$ for Model 6, $p < 0.001$ for both). Responsiveness also differs between the two cells in each market. In the 2x2 market, buyers significantly under–respond to price differences when both are informed, but when only one buyer is informed – so that calculating the predicted probabilities is trivial, even for the participants – responsiveness is either weakly significantly higher (Model 6) or the increase just misses significance (Model 5). We see the opposite relationship in the 2x3 market, with responsiveness weakly significantly *lower* when one buyer is uninformed than when all buyers are informed.

The bottom half of the table adds additional detail. As with sellers (Table 5), buyer risk aversion correlates in the expected way with behavior in the market. Here, the most risk–averse buyers are significantly less price–responsive than the least risk–averse ones; since (except in the 2x2–1 cell) visiting the lower–priced seller entails a higher profit given trading but a lower probability of trading, risk aversion should indeed lower the likelihood of choosing that seller. Also, buyers' views of fairness have the expected association with their market behavior (see Section 3.3); buyers who report a fair price of 8 (roughly the first quartile) are significantly more price–elastic than those who report a fair price of 13 (roughly the third quartile). Finally, a comparison between Models 6 and 7 suggests that all of the effects described here are present in early rounds as well as overall, though in some cases they become less significant due to the larger standard errors arising from smaller sample sizes.

Result 3 *In all four markets, buyers exhibit lower price elasticity than predicted. This under–responsiveness varies systematically with the cell, and is least severe in the 2x2–1 cell and most severe in the 2x3–1 cell.*

Result 4 *Buyers' price–responsiveness varies with their risk attitudes and views of fairness. In particular, they are more price–responsive if they are less risk–averse and if they view lower prices as being fair.*

Note that the last part of Result 4 is consistent with Proposition 1.

4.4 Seller prices revisited

Since the predictions in Table 1 are based on equilibrium behavior by informed buyers (with standard preferences), and since we have just seen that buyers are actually insufficiently price–responsive relative to that benchmark, it may be that sellers' prices, while not consistent with subgame perfect equilibrium, do represent Nash equilibrium choices *given buyers' actual behavior*. We now examine this possibility.

In order to compute Nash equilibrium seller prices, we need to specify buyers' continuation behavior. The assumption we make is that buyers are globally β times as price–responsive as the theoretical prediction under standard preferences. That is, when the symmetric Nash equilibrium in buyer strategies has each informed buyer choosing Seller 1 with probability $q(p_1, p_2)$, where p_1 and p_2 are the sellers' price choices, the actual probability of an informed buyer choosing Seller 1 is

$$\hat{q}(p_1, p_2) = \frac{1-\beta}{2} + \beta \cdot q(p_1, p_2). \tag{3}$$

So, when sellers ought to be visited with equal probability ($q = \frac{1}{2}$), this is what happens ($\hat{q}$ is also $\frac{1}{2}$). Also, $\frac{\partial \hat{q}}{\partial p_i} = \beta \cdot \frac{\partial q}{\partial p_i}$ for $i = 1, 2$: price–responsiveness is β times the predicted level.

While simply inserting a constant multiplier into the visit–probability function may seem ad hoc, this is not necessarily the case. In particular, if all buyers have utility function $u(x) = x^\beta$ for $\beta \in (0, 1)$ (constant relative risk–aversion, or CRRA), then they will be exactly β times as price–responsive in equilibrium as they would be under risk–neutrality. In Appendix A.3, we derive the theoretical predictions for the 2x3–0 and 2x3–1 cases under general risk–averse utility functions; substituting a CRRA utility function for the general function (e.g., in (8) in the appendix) demonstrates the result we mention here. Also, in Appendix A.5 we directly compute the equilibrium in seller prices for the 2x2–0 case under the assumption in (3).

Figure 7 depicts how the Nash equilibrium prices for the four cells of our experiment depend on β over a range of values from complete unresponsiveness ($\beta = 0$) to full responsiveness ($\beta = 1$), and even over–responsiveness ($\beta > 1$). (Note that over–responsiveness is not possible in the 2x2–1 case, since full responsiveness means choosing the lower–price seller with certainty.) Within any cell, higher price responsiveness

Figure 7: Nash equilibrium price choices, given informed–buyer price responsiveness of β times predicted responsiveness

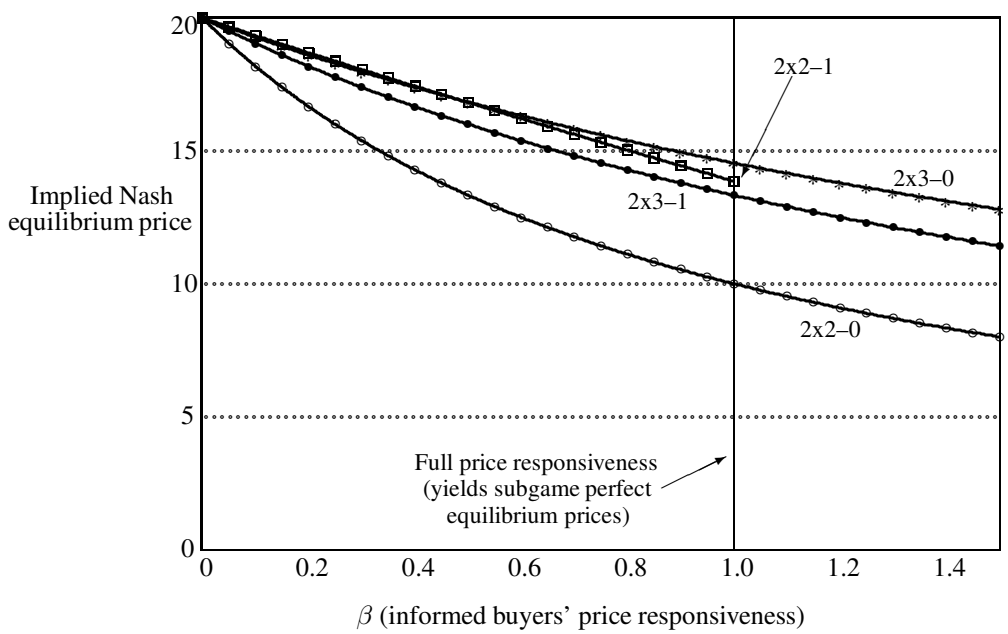

implies lower prices. Also, within either market (2x2 or 2x3), the ordinal relationship between zero and one uninformed buyers is the same over the range of β pictured. This last fact means that our 2x3–market result cannot be explained as arising from equilibrium seller behavior given buyer under–responsiveness that is constant across treatments (e.g., if under–responsiveness were due to risk aversion).

To get the predicted Nash equilibrium prices, it remains only to specify values for β. These we estimate from the data: specifically, using the individual–cell estimates from Model 6 in Table 6. The resulting prices (point estimates and 95% confidence intervals) are shown in Table 7. Nash–equilibrium prices are higher

Table 7: Predicted and observed posted prices

Cell	Subgame perfect	Observed price		Nash equil. price, given buyer responsiveness	
	equil. price	all rounds	last 5 rounds	point estimate	95% CI
2x2–0	10.00	9.69	9.85	11.79	(10.82, 12.76)
2x2–1	13.86	12.30	12.81	14.73	(14.22, 15.24)
2x3–0	14.55	11.86	12.98	15.93	(15.37, 16.50)
2x3–1	13.33	13.19	14.19	15.67	(15.08, 16.26)

than the corresponding subgame–perfect–equilibrium price, reflecting buyer under–responsiveness. In the 2x2 market, the impact of using actual buyer price–responsiveness instead of assuming full responsiveness is minor. The higher responsiveness in the 2x2–1 cell compared with the 2x2–0 cell means that the Nash equilibrium price is about a dollar higher than the subgame perfect equilibrium price in the former, versus almost two dollars higher in the latter. However, while this reduces the size of the predicted treatment effect (from 3.86 to 2.94), it does not change its sign, and indeed these point estimates are significantly different from each other ($p < 0.001$), as the 95% confidence intervals would suggest. So, accounting for buyers' actual level of price–responsiveness rather than assuming the equilibrium level has no material effect on the theoretical predictions for sellers in the 2x2 market.

By contrast, in the 2x3 market, the higher price responsiveness in the 2x3–0 cell relative to the 2x3–1 cell is sufficient to reduce the predicted treatment effect nearly to zero (from 1.22 to 0.26), and as the nearly identical confidence intervals suggest, there is no significant difference between the point estimates ($p \approx 0.52$). Thus, given buyer behavior, there is essentially no monetary incentive for sellers to post lower prices in the 2x3–1 cell compared to the 2x3–0 cell as predicted by the subgame perfect equilibrium.

Result 5 *In the 2x2 market, differences in buyer price responsiveness between the 2x2–0 and 2x2–1 cells are small enough to leave the predicted treatment effect (due to competitive pressures alone) intact. In the 2x3 market, corresponding differences are large enough to reduce the predicted treatment effect almost to zero.*

Next, we perform similar probit estimations to those in Table 6, but separately for each matching group and cell, each yielding corresponding estimates of β and thence Nash equilibrium seller price. These are plotted against the associated mean observed seller prices in Figure 8. While there is substantial heterogeneity across groups, sellers on average under–price relative to the Nash equilibrium price.

Table 8 shows additional under–pricing results based on Tobit regressions on the individual–level data, with the level of under–pricing (the posted price minus the Nash equilibrium price for that cell and matching group) as the dependent variable. Explanatory variables include dummies for the 2x2–1, 2x3–0 and 2x3–1 cells, the round number, a second–half dummy, a constant term, and a new variable called "fair price

Figure 8: Scatter–plot of Nash equilibrium seller prices (based on actual buyer behavior) and actual seller posted prices, by individual matching group

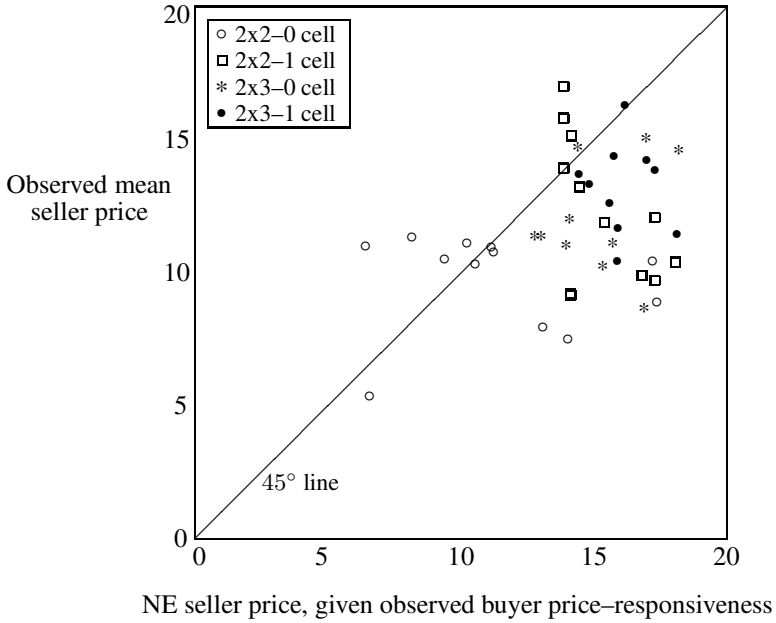

deviation", which is the seller's reported fair price for the current round minus the Nash equilibrium price. Along with marginal effects and standard errors, we report estimated levels of under–pricing in each cell, and the corresponding 95% confidence intervals. One set of estimates is based on all seller data, one is based on only the first five rounds of each cell, and a third is based on the last five rounds of each cell.

Sellers under–price in all four cells; it decreases over time but remains substantial even in the last five rounds. As with buyer price–responsiveness, there are systematic differences across cells. Most importantly, under–pricing is significantly more severe in the 2x3–0 cell than in the 2x3–1 cell. Finally, we again find a correlation between reported fair prices and actual price choices, not only over all rounds but importantly also in the first five rounds on their own (suggesting that the correlation is not due simply to sellers calculating the average price for the session and reporting that as their fair price).

Result 6 *Sellers under–price relative to the Nash equilibrium price given buyers' actual price–responsiveness. This under–pricing varies systematically with the cell, and overall is most severe in the 2x3–0 cell.*

Result 7 *Sellers under–price more severely the lower is their reported fair price.*

Note that the last part of Result 7 is consistent with Proposition 2.

4.5 How do fair prices vary across treatments?

The last three sections have shown evidence consistent with our Propositions 1 and 2, based on estimations that either were carried out separately for each cell (in the case of buyers) or included controls for the cell

Table 8: Tobit results – estimated average marginal effects (unless specified) and standard errors

Dependent variable: price deviation (posted price minus Nash equilibrium price given buyer price–responsiveness)				Associated level of under–pricing (point estimate and 95% CI)			
	Rounds 1–5	Rounds 16–20	All rounds		Rounds 1–5	Rounds 16–20	All rounds
				2x2–0	–2.529	–2.398	–2.515
					(–3.179,–1.878)	(–3.060,–1.736)	(–3.099,–1.931)
2x2–1	-0.799^{***}	0.546^{***}	0.125	2x2–1	–3.329	–1.852	–2.390
	(0.238)	(0.185)	(0.107)		(–3.973,–2.686)	(–2.509,–1.195)	(–2.972,–1.808)
2x3–0	-2.253^{***}	0.095	-0.967^{***}	2x3–0	–4.786	–2.303	–3.482
	(0.476)	(0.487)	(0.431)		(–5.455,–4.116)	(–2.988,–1.617)	(–4.092,–2.873)
2x3–1	-1.841^{***}	0.694	–0.195	2x3–1	–4.371	–1.704	–2.710
	(0.478)	(0.488)	(0.431)		(–5.041,–3.701)	(–2.390,–1.018)	(–3.319,–2.100)
Joint significance	$p < 0.001$	$p < 0.001$	$p < 0.001$				
Significance, 2x3–0 vs. 2x3–1	$p \approx 0.066$	$p < 0.001$	$p < 0.001$				
Round number	0.202^{***}	0.127^{***}	0.105^{***}				
	(0.055)	(0.042)	(0.006)				
Second half	1.482^{***}	0.999^{***}	1.231^{***}				
	(0.159)	(0.122)	(0.070)				
Seller fair price deviation	0.582^{***}	0.650^{***}	0.620^{***}				
(fair price minus NE price)	(0.043)	(0.035)	(0.021)				
N	920	920	3680				
$\|ln(L)\|$	2186.59	1968.33	8063.83				

* (**,***): *Marginal effect significantly different from zero at the 10% (5%, 1%) level.*

(for sellers), and are thus *within–cell* relationships. However, if there are systematic differences in these reported fair prices *across cells*, these might provide empirical support for the fairness explanation for the treatment effects we have observed.

Figure 9 displays scatter–plots of individual participants' reported fair prices with zero and one uninformed buyer, with buyers displayed as dark circles and sellers as light circles/rings. (Recall that each participant played 20 rounds with zero uninformed buyers and 20 rounds with one uninformed buyer, and reported separate fair prices for each.) Also shown are the corresponding means for all buyers and sellers in each cell, and the subgame–perfect–equilibrium prices under standard preferences. While the modal fair price for both buyers and sellers is the equal–split price of 10, there is substantial variation within and across cells. The within–cell variation suggests that participants are heterogeneous in their fairness views. The across–cell variation suggests that these views are not completely home–grown; rather, they seem partly to reflect participants' perception of differing strategic environments across cells. In both markets, both buyers and sellers report fair prices that are significantly higher with one uninformed buyer than with zero uninformed buyers (Wilcoxon signed–ranks tests, $p < 0.01$ for sellers in either market, $p \approx 0.018$ for buyers in

Figure 9: Participants' reported fair prices and subgame perfect equilibrium prices under standard preferences *(area of a dark circle is proportional to the number of buyers, area of a light circle/ring is proportional to the number of sellers it represents)*

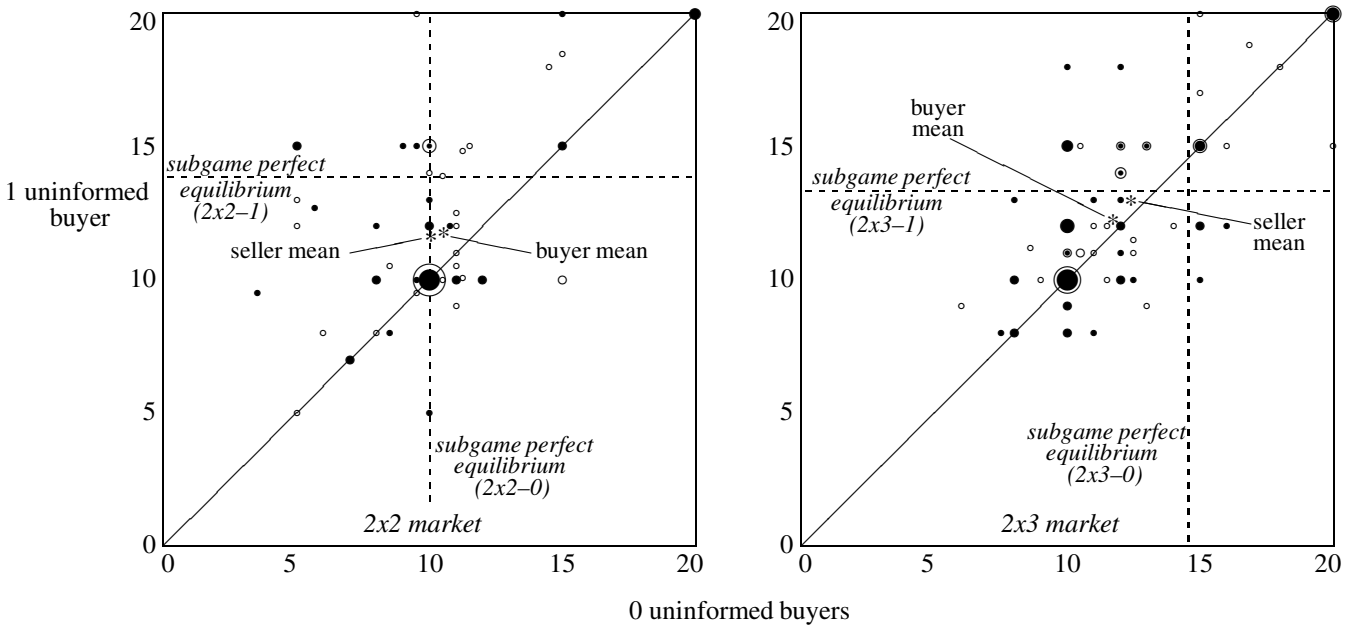

the 2x2 market, $p \approx 0.062$ for buyers in the 2x3 market).[15] Most striking, the higher fair price in the 2x3–1 cell compared to the 2x3–0 cell is the reverse of the equilibrium prediction. If we define a participant–specific "fair price deviation" measure as the difference between the subgame perfect equilibrium price in a cell and the participant's reported fair price for that cell, this deviation is significantly *larger* (i.e., fair prices under–state the equilibrium price *more*) in the 2x2–1 cell than in the 2x2–0 cell, while it is significantly *smaller* in the 2x3–1 cell than in the 2x3–0 cell (Wilcoxon signed–ranks tests, $p < 0.001$ in both markets and for both buyers and sellers).

Result 8 *Both buyers and sellers report* higher *fair prices in the 2x2–1 cell than in the 2x2–0 cell. Both buyers and sellers report* higher *fair prices in the 2x3–1 cell than in the 2x3–0 cell.*

Summing up, the differences between the observed treatment effects and the theoretical predictions under standard preferences are consistent with the implications of fairness. In the 2x2 market, competitive pressures alone predict a price increase when we move from zero to one uninformed buyers. While the associated difference in fair prices is smaller than the predicted price difference, it has the same sign. Since we have seen that higher fair prices are associated with higher actual prices (both directly in seller choices,

[15] Curiously, we find only weak evidence of self–serving bias. In the 2x2 market, buyers actually report *higher* fair prices than sellers, though these differences are not significant (robust rank–order tests, $p > 0.20$ for both 2x2–0 and 2x2–1). In the 2x3 market, sellers do report at least weakly significantly higher fair prices than buyers ($p \approx 0.043$ for 2x3–0, $p \approx 0.091$ for 2x3–1), though the magnitude of the difference, roughly 70 cents in both cases, is small.

and indirectly through the incentive for higher prices due to reduced buyer price–responsiveness), we would therefore expect fairness to reinforce the effects of competitive pressures, meaning higher prices in the 2x2–1 cell than in the 2x2–0 cell, which is indeed what we observe. In the 2x3 market, competitive pressures alone predict a price decrease when we move from zero to one uninformed buyers. However, the difference in fair prices is in *the opposite direction*, which implies higher prices in the 2x3–1 cell than in the 2x3–0 cell. Thus what we observe is exactly consistent with the "fairness" alternative to Hypothesis 2: the effect of fairness has the opposite sign of competition's effect, and is apparently strong enough to flip the direction of the overall effect from that of the standard–theory prediction.

4.6 Discussion of our results as compared with Helland et al.'s (2016)

Our finding of lower prices in the 2x3–0 cell than in the 2x3–1 cell contrasts with another current study (Helland et al. 2016), whose experiment implements a similar setting to ours, and who report strong support for the theoretical predictions based on standard preferences. In particular, they find at least weakly significant support for the counter–intuitive prediction of *higher* prices in their 2x3–0 cell than in their 2x3–1 cell.[16] Their experimental procedures differ in many ways from ours; however, both sets of procedures arguably conform to standard experimental–economics methodology, so neither is manifestly "correct" or "wrong".[17]

The most plausibly consequential difference between the two experiments is that their participants play 50 rounds of one treatment, while ours play 20 rounds each of two treatments. This could potentially explain the differing results we find, since one might expect that giving participants more experience in the same setting ought to make them play more in line with theoretical predictions. Indeed, even though Helland et al. observe essentially no time trend in prices after the 20th round (see their Figure 2) – suggesting their results would have been the same if they had stopped at that point – prices in our 2x3–0 and 2x3–1 cells do not appear to have stabilized at the 20th round, and while the *difference* between these prices does appear to remain fairly stationary after the first few rounds, it is always possible that this could have changed in later

[16] Based on the one–sided non–parametric tests they report, the difference they find in posted prices is statistically insignificant and the one in transaction prices is significant at the 10% level. Based on regressions, the difference is significant at the 10% level for posted prices and the 5% level for transaction prices. Given that non–parametric tests tend to err on the side of conservatism and regressions on the side of liberalism, these results probably indicate borderline significance.

[17] To our knowledge, here are the differences between our original experiment and Helland et al.'s (2016), other than those emphasized in the main text of this section. Helland et al. use a between–participant design while we use a within–participant design. Their uninformed buyers are automated while ours are human participants. They assign persistent labels to sellers (which allows collusion and coordination) while we do not. Their matching groups comprise three individual markets; ours comprise either two or three, and we control for the size in our seller regressions. They paid participants the sum of payoffs in all rounds multiplied by a conversion factor, while we pay the actual payoffs from a subset of rounds. They use only 2x3 markets but also vary whether a capacity constraint exists, while we always have capacity constraints but vary the market size. Their sessions were conducted in Norway and Germany; ours were conducted in Australia. Their buyer valuation was 100 experimental currency units, while ours was 20 dollars. They do not elicit risk attitudes, demographic information, or perceived fair prices; we elicit all of these. Finally, they use one–sided non–parametric tests while we use more conservative two–sided tests; this obviously cannot change qualitative treatment effects, but does make their significance results appear stronger than they are.

rounds had they been played.

To address this possibility, we report the results of a small follow–up experiment in which participants play 50 rounds under a single set of parameters (as in Helland et al.): either the 2x3 market with zero uninformed buyers or the 2x3 market with one uninformed buyer. Otherwise, experimental procedures parallel those in our original experiment (to keep the stakes per round comparable, we paid according to 5 rather than 4 rounds). We conducted 6 sessions – 3 each of 2x3–0 and 2x3–1 – with 20 participants in each session, none of whom had participated in our original experiment, and without partitioning the session into smaller matching groups. (So, each matching group comprised 4 markets.)

Two key results are apparent in Figure 10, which shows the time series of posted and transaction prices in both treatments. First, the counter–theoretic but intuitive (and fairness–consistent) result of *higher* prices

Figure 10: Time series of posted and transaction prices in 2x3–0 and 2x3–1 treatments of new experiment

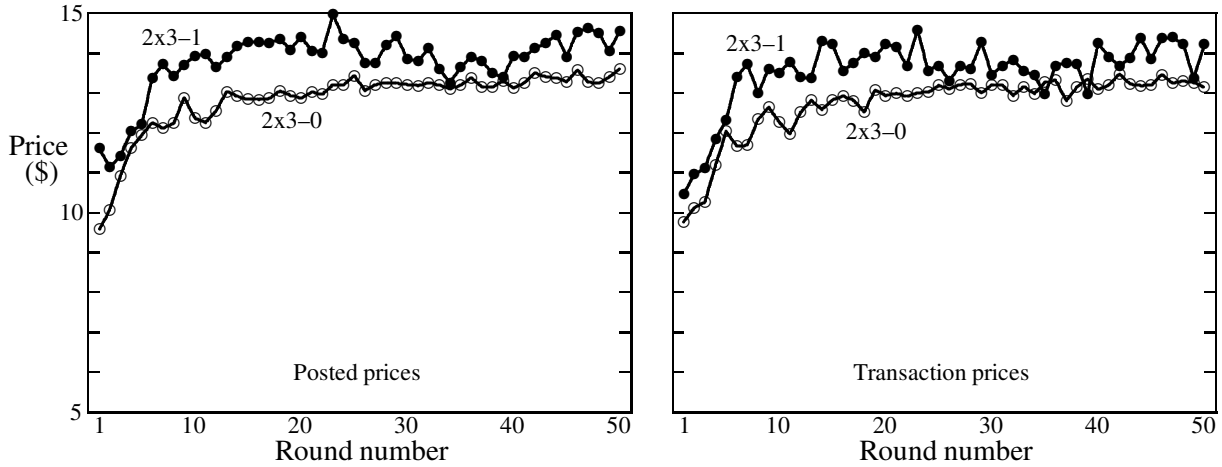

in 2x3–1 compared to 2x3–0 is present in the new experiment as it was in the original one.[18] Second, while prices continue rising after round 20, these rises are small in magnitude and not substantially different between the treatments. (Indeed, in the last ten rounds the treatment effect appears to get larger, not smaller.) Thus it is very likely that the differences between our original experimental results and those of Helland et al. are not due to the number of rounds played (or indeed within– versus between–participant design).

With these potential explanations ruled out, and since none of the other differences between our experiments (see footnote 17) seems like a plausible explanation, we are left merely to speculate. One possibility, as always, is differences in the participant pool. If Helland et al.'s participants did not care about fair prices, or if their fair–price concerns were different from our participants' (e.g., if the same price was viewed as equally fair across all treatments), that could easily account for the difference between our results and theirs. Alternatively, Helland et al.'s between–participants design allows for uncontrolled variation in other char-

[18] As in the original experiment, there is some heterogeneity across sessions. In the three 2x3–0 sessions, posted prices averaged 12.07, 12.51 and 13.95 and transaction prices averaged 12.00, 12.27 and 13.85, while in the three 2x3–1 sessions, posted prices averaged 13.03, 13.74 and 14.61 and transaction prices averaged 12.72, 13.41 and 14.48.

acteristics such as risk attitudes: if their buyers are less risk averse, or sellers more risk averse, in their 2x3–1 treatment than in 2x3–0, that could explain their result.[19] There are also other potential explanations – perhaps there is some difference in procedures that we have failed to notice – but again, we can do little beyond speculating.

5 Conclusion

While the results in our 2x2 market are consistent with standard theory, the same is not true for our 2x3 market. There, prices are *higher* with one uninformed buyer (2x3–1) than with zero (2x3–0), rather than *lower* as predicted. This result cannot be due to risk aversion, since (a) we vary the number of uninformed buyers *within–participants* to eliminate differences in risk attitudes between the 2x3–0 and 2x3–1 cells, and (b) our analyses in Appendices A.3 and A.4 show that the sign of the predicted difference between these two cells is robust to allowing either identical or heterogeneous risk aversion, as long as risk attitudes do not change with the number of uninformed buyers. It cannot be due to order effects, as it is confirmed both in robustness checks using the data from the original experiment (footnotes 11 and 12) and in a follow–up experiment using a between–participants design (Section 4.6). The follow–up experiment also shows that the result is not due to allowing insufficient time for play to converge.

Instead, we find evidence suggesting that market outcomes are influenced by fairness perceptions. Reported fair prices vary widely within each cell and in systematic ways across cells. Like observed prices but not theoretically–predicted prices, fair prices in the 2x3 market are significantly higher with one uninformed buyer than with zero. (In the 2x2 market, fair prices show the same relationship as both observed prices and predicted prices.) Since fair prices were elicited at the end of the experimental session, after all of the market decisions were made, we cannot prove causality. However, a compelling (though still circumstantial) argument is made by within–cell analysis, which shows that participants' reported fair prices correlate with their behavior – even in early rounds where predicted and/or observed treatment effects would not have become apparent. Buyers reporting lower fair prices are more price–responsive (thus incentivizing sellers to lower their prices), while sellers reporting lower fair prices charge lower prices, even after accounting for the effect of buyers' actual price–responsiveness. Both of these patterns of behavior are consistent with rationality if standard preferences are augmented by a distaste for prices away from a "fair price" reference point, as discussed in Section 3.3. Thus, the counter–intuitive nature of the standard theoretical prediction for the 2x3 market may be, in a sense, exactly what leads to its undoing: participants quickly form views

[19] Since Helland et al. do not collect information about participant characteristics, we can only conjecture about whether their participant pool is different from ours. A hint that their participants may be unusual is given by buyer visit choices. The under–responsiveness to price in our original experiment is similar to that seen in other directed–search experiments (Cason and Noussair 2007, Anbarci and Feltovich 2013). Our follow–up experiment is similar: a probit like Model 6 yields marginals of 0.738 and 0.570 for the 2x3–0 and 2x3–1 cells, both significantly less than one. By contrast, Helland et al. report approximately full responsiveness (from their Result 6, p. 12: "the average probability of buying at a specific seller is almost identical to the average predicted probability given prices"), though their Figure 4 suggests that this holds in their 2x3–1 cell better than in their 2x3–0 cell.

about "fair" prices – perhaps based entirely on intuition that more information should reduce prices – and their resulting behavior suggests these views are self–fulfilling.

Our results are in keeping with previous fairness results in market experiments. Like Kachelmeier et al. (1991a), we find individual–level correlations between perceived fair prices and market behavior, and like their study and several follow–ups (e.g., Kachelmeier et al. 1991b, Franciosi et al. 1995, Ruffle 2000, Bartling et al. 2014), we find that the success of theoretical predictions is influenced by how fair the predicted outcomes are viewed to be. So, part of our contribution is simply the added evidence of the robustness of fairness effects in markets. However, our study is also novel in several ways. One is in the milieu we use: markets with directed search. Our procedures make this a severe test of the effects of fairness on buyer behavior: unlike Kachelmeier et al. (1991a; 1991b), whose buyers could post bids; Franciosi et al. (1995), Ruffle (2000), and Tyran and Engelmann (2005), whose buyers could refuse to buy at all; and Bartling et al. (2014), whose sellers could engage in punishment; our buyers have no recourse when facing unfair prices except to visit a different seller. Even limited to such a weak punishment mechanism, buyers utilize it and its effects are apparent. Our results are also novel in that they are persistent (with no apparent decrease over 50 rounds in our follow–up experiment), in contrast to previous results (Kachelmeier et al. 1991a; 1991b; Franciosi et al. 1995) where the treatment effect is in speed of convergence to the same equilibrium, and thus transitory. Finally, fairness perceptions in our experiment actually reverse a theoretically predicted treatment effect (between the 2x3–0 and 2x3–1 cells), which we have not seen elsewhere.

We expect that our results have clear implications for market settings outside the lab. While external validity is always a concern when drawing conclusions based on lab experiments, we are relatively confident about our findings for buyer behavior, since our participants have plenty of experience as buyers in many settings, and – as undergraduates at a large public university – they are fairly representative of the overall population except in age, which seems not to be a significant driver of behavior here. We are less confident that our findings for sellers will generalize, since undergraduates are less representative and less experienced as sellers than as buyers. However, as we saw in Section 4.4, even if sellers play exactly as Nash equilibrium implies given buyers' behavior, buyers' deviations from equilibrium behavior on their own are sufficient to all but eliminate the predicted treatment effect between our 2x3–0 and 2x3–1 cells. Since this pair of cells represents the largest counter–intuitive price difference that is possible within the model, even for more general market sizes, and even if "real" sellers were to behave in line with the theory, it seems unlikely that we would see a case where more informed buyers led to higher prices.

Our results also have implications for policy. Lester (2011) provides some motivation for his paper by noting that many US states have legislated to increase price transparency in health–care markets. His theoretical results suggest that such policies may have perverse effects, since making prices more transparent may lead to price increases rather than decreases. Our empirical results suggest that such perverse effects may be less likely than the theory implies, tilting the balance in favor of price–transparency policies.

References

Albæk S, Møllgaard P, Overgaard PB (1997) Government–assisted oligopoly coordination? A concrete case. *J. Indust. Econom.* 45:429–443.

Anbarci N, Dutu R, Feltovich N (2015) Inflation tax in the lab: a theoretical and experimental study of competitive search equilibrium with inflation. *J. Econom. Dynamics and Control* 61:17–33.

Anbarci N, Feltovich N (2013) Directed search, coordination failure and seller profits: An experimental comparison of posted pricing with single and multiple prices. *Internat. Econom. Rev.* 54:873–884.

Babcock L, Loewenstein G, Issacharoff S, Camerer C (1995) Biased judgments of fairness in bargaining. *Amer. Econom. Rev.* 85:1337–1343.

Bartling B, Greider M, Zehnder C (2014) Let the market decide: an experimental study of competition and fairness. CESifo working paper 4831.

Burdett K, Shi S, Wright R (2001) Pricing and matching with frictions. *J. Political Economy* 109:1060–1085.

Campbell CL, Heinrich D, Schoenmüller V (2015) Consumers' reaction to fair trade motivated price increases. *J. Retailing and Consumer Services* 24:79–84.

Campbell, MC (1999) Perceptions of price unfairness: antecedents and consequences. *J. Marketing Research* 36:187–199.

Cason TN, Gangadharan L, Nikiforakis N (2011) Can real–effort investments inhibit the convergence of experimental markets?" *Internat. J. Indust. Organ.* 29:97–103.

Cason TN, Noussair C (2007) A market with frictions in the matching process: an experimental study. *Internat. Econom. Rev.* 48:665–691.

Dickson PR, Kalapurakal R (1994) The use and perceived fairness of price–setting rules in the bulk electricity market. *J. Econom. Psych.* 15:427–448.

Eckel CC, Grossman PJ (2008) Forecasting risk attitudes: an experimental study of actual and forecast risk attitudes of women and men. *J. Econom. Behav. Organ.* 68:1–17.

Fehr E, Kirchsteiger G, Riedl A (1993) Does fairness prevent market clearing? An experimental investigation. *Quart. J. Econom.* 108:437–459.

Feltovich N (2005) Critical values for the robust rank–order test. *Communications in Statistics – Simulation and Computation* 34:525–547.

Fischbacher U (2007) z–Tree: Zurich toolbox for ready–made economic experiments. *Experiment. Econom.* 10:171–178.

Franciosi R, Kujal P, Michelitsch R, Smith V, Deng G (1995) Fairness: effect on temporary and equilibrium prices in posted–offer markets. *Econom. J.* 105:938–950.

Helland L, Moen E, Preugschat E (2016) Information and coordination frictions in experimental posted offer markets. Working paper.

Kachelmeier SJ, Limberg ST, Schadewald MS (1991a) A laboratory market examination of the consumer price response to information about producers' costs and profits. *Accounting Rev.* 66:694–717.

Kachelmeier SJ, Limberg ST, Schadewald MS (1991b) Fairness in markets. *J. Econom. Psych.* 12:447–464.

Kahneman D, Knetsch JL, Thaler R (1986) Fairness as a constraint on profit seeking: entitlements in the market. *Amer. Econom. Rev.* 76:728–741.

Kloosterman A (2016) Directed search with heterogeneous firms: an experimental study. *Experiment. Econom.* 19:51–66.

Korth C (2009) Price rigidity in an experimental market. Korth C, ed. *Lecture Notes in Econom. and Mathematical Systems 627: Fairness in Bargaining and Markets* (Springer, Berlin), 139–167.

Lester B (2011) Information and prices with capacity constraints. *Amer. Econom. Rev.* 101:1591–1600.

McKelvey RD, McLennan AM, Turocy TL (2014) Gambit: software tools for game theory. Version 13.1.2, `www.gambit-project.org`.

Piron R, Fernandez L (1995) Are fairness constraints on profit–seeking important?" *J. Econom. Psych.* 16:73–96.

Renner E, Tyran J-R (2004) Price rigidity in customer markets. *J. Econom. Behav. Organ.* 55:575–593.

Roth AE (1995) Bargaining experiments. Kagel JH, Roth AE, eds., *Handbook of Experimental Economics* (Princeton University Press, Princeton, NJ), 253–348.

Ruffle BJ (2000) Some factors affecting demand withholding in posted–offer markets. *Econom. Theory* 16:529–522.

Salop S, Stiglitz JE (1977) Bargains and ripoffs: a model of monopolistically competitive price dispersion. *Rev. Econom. Studies* 44:493-510.

Seligman DA, Schwartz B (1997) Domain specificity of fairness judgments in economic transactions. *J. Econom. Psych.* 18:579–604.

Siegel S, Castellan NJ Jr (1988) *Nonparametric Statistics for the Behavioral Sciences* (McGraw–Hill, New York).

Thaler RH (1985) Mental accounting and consumer choice. *Marketing Science* 4:199–214.

Tian Y, Zhou H (2015) From bottom line to consumers' mind: the framing effects of accounting information. *Accounting, Organ. and Soc.* 43:56–66.

Tyran J-R, Engelmann D (2005) To buy or not to buy? An experimental study of consumer boycotts in retail markets. *Economica* 72:1–16.

Varian HR (1980) A model of sales. *Amer. Econom. Rev.* 70:651–659.

Xia L, Monroe KB, Cox JL (2004) The price is unfair! A conceptual framework of price fairness perceptions. *J. Marketing* 68:1–15.

A Appendix

A.1 Solution to the case of one informed buyer (standard preferences)

Suppose there are $m \geq 2$ sellers and $n \geq 2$ buyers, with exactly 1 informed buyer and the remaining $n-1$ buyers uninformed. Let $\bar{U}$ be buyers' common valuation for the good (in the experiment, $\bar{U} = 20$).

Proposition 3 *Define*

$$p^* = \bar{U}\left[1 - \left(1 - \frac{1}{m}\right)^{n-1}\right]. \tag{4}$$

Then there exists a unique symmetric subgame perfect equilibrium. In this equilibrium, each seller's price is drawn from a distribution with support $[p^*, \bar{U}]$ and no mass points. For p in this interval, the cumulative distribution function (c.d.f.) F is given by

$$F(p) = 1 - \left[\frac{p^*(\bar{U} - p)}{p(\bar{U} - p^*)}\right]^{\frac{1}{m-1}}. \tag{5}$$

Proof: Buyer behavior is straightforward: each of the u uninformed buyers mixes uniformly over sellers, and the one informed buyer chooses the lowest–priced seller with certainty, or mixes among them if there is more than one lowest–priced seller. So, a seller choosing price $p \leq \bar{U}$ will sell for sure if it is the lowest price (since the informed buyer visits); while if another seller has a strictly lower price, she will sell with probability $1 - \left(1 - \frac{1}{m}\right)^{n-1}$, the probability of at least one uninformed buyer visiting. Finally, if she is one of exactly k lowest–priced sellers, she will sell with probability $1 - \left(1 - \frac{1}{k}\right)\left(1 - \frac{1}{m}\right)^{n-1}$; note that this is strictly larger than $1 - \left(1 - \frac{1}{m}\right)^{n-1}$ for any $k \in \{1, 2, ..., n\}$, and strictly smaller than one as long as $k > 1$ (i.e., the seller is tied with at least one other seller).

Let F be the common c.d.f. of the sellers. F cannot have a mass point at any price $p \in [p^*, \bar{U}]$, since then a seller could improve her profit by shifting some mass from p slightly to the left (removing the chance of a tie and hence discontinuously raising the probability of selling, with only a second–order decrease in the selling price).[20] Also, F cannot have a "gap" – an interval (p_0, p_1) with zero density throughout but positive density below and above – since then choosing price p_1 would earn a strictly higher profit than prices just below p_0, as it would increase the selling price with only a second–order decrease in the probability of selling.

Let $\hat{p} = sup\{p|F(p) = 0\}$ and $\tilde{p} = min\{p|F(p) = 1\}$ (the latter exists because of right–continuity of c.d.f.'s). Then choosing $\tilde{p}$ yields profit $\pi(\tilde{p}) = \tilde{p}\left[1 - \left(1 - \frac{1}{m}\right)^{n-1}\right]$. But choosing price $\bar{U}$ yields profit $\pi(\bar{U}) = \bar{U}\left[1 - \left(1 - \frac{1}{m}\right)^{n-1}\right]$, so we cannot have $\tilde{p} < \bar{U}$. Also, any price above $\bar{U}$ yields zero profit, so we cannot have $\tilde{p} > \bar{U}$. Hence we must have $\tilde{p} = \bar{U}$.

Then, choosing $\hat{p}$ yields profit $\pi(\hat{p}) = \hat{p}$, which must be equal to $\pi(\bar{U})$; that is, $\hat{p} = p^*$ as defined in (4). So, p^* and $\bar{U}$ are the left and right endpoints of the smallest closed interval containing the support of

[20] The arguments for this "no mass point" assertion and the upcoming "no gap" assertion are based on arguments used by Varian (1980, pp. 653–654) in his model of temporal price dispersion.

the price distribution. Since we have already noted there are no "gaps" in this support, this means that the support of F is exactly $[p^*, \bar{U}]$.

Finally, let $p \in [p^*, \bar{U}]$. Since F has no mass points, the probability of another seller choosing exactly p is zero, so we can ignore the possibility of ties for the lowest price. Then p is the lowest price with probability $[1 - F(p)]^{m-1}$, in which case it sells for sure; otherwise, another seller's price is strictly lower, in which case it sells with probability $1 - \left(1 - \frac{1}{m}\right)^{n-1} = p^*/\bar{U}$. So, the resulting expected profit is given by

$$\pi(p) = p \cdot [1 - F(p)]^{m-1} + p \cdot \frac{p^*}{\bar{U}} \left\{1 - [1 - F(p)]^{m-1}\right\}$$

Since there is positive density at p, it must yield profit equal to that of p^*, so

$$p^* = p \cdot [1 - F(p)]^{m-1} + p \cdot \frac{p^*}{\bar{U}} \left\{1 - [1 - F(p)]^{m-1}\right\}.$$

Solving for $F(p)$ yields (5). □

Note that for the special case of the 2x2–1 market, (5) simplifies to $F(p) = 2 - 20/p$, so the corresponding probability density function (p.d.f.) is $f(p) = 20/p^2$ with support [10, 20] and expected value $E(p) = 20 \cdot ln(2) \approx 13.86$.

A.2 Proofs of Propositions 1 and 2 (2x2–0 case)

Proposition 1

Let $\bar{U}$ be buyers' common valuation for the good (in the experiment, $\bar{U} = 20$). Suppose the sellers post prices p_1 and p_2, not both equal to $\bar{U}$, and with $p_1 \leq p_2$ without loss of generality. Define $c_i = c(p_i - p^*)$.

(i) Then in the symmetric equilibrium of the game played between buyers, the probability a given buyer visits Seller 1 is given by

$$q^* = \frac{\bar{U} - 2p_1 + p_2 - 2c_1 + 2c_2}{2\bar{U} - p_1 - p_2}, \tag{6}$$

when this is between 0 and 1.[21] Note that when $c_1 = c_2 = 0$, this expression reduces to (2) in Section 4.3.

Now,

$$\frac{\partial q^*}{\partial p^*} = \frac{2}{2\bar{U} - p_1 - p_2} \left(\frac{\partial c_2}{\partial p^*} - \frac{\partial c_1}{\partial p^*}\right),$$

and note that the expression outside the parentheses is positive, so that $\frac{\partial q^*}{\partial p^*}$ has the same sign as $\left(\frac{\partial c_2}{\partial p^*} - \frac{\partial c_1}{\partial p^*}\right)$. Since $p_1 \leq p_2$, this means that $\frac{\partial q^*}{\partial p^*} < 0$ if $c(\cdot)$ is convex (which is true by assumption), and also that q^* is the probability of visiting the lower–priced seller. Thus, as p^* increases, the equilibrium probability of visiting the lower–price seller decreases. □

[21] Since at least one price is less than $\bar{U}$, q^* exists. If the right–hand side of (6) is less than zero (more than one), buyers visit Seller 2 (Seller 1) with certainty.

(ii) Suppose now that p^* varies across buyers, with a distribution over $[0, \bar{U}]$ that is common knowledge. Suppose that each buyer's value of p^* is his private information. Then a given buyer's expected utility from visiting Seller 1 and Seller 2 are

$$q\left[\frac{\bar{U}-p_1}{2}\right] + (1-q)[\bar{U}-p_1] - c(p_1-p^*) = \left(1-\frac{q}{2}\right)[\bar{U}-p_1] - c(p_1-p^*)$$

and

$$q[\bar{U}-p_2] + (1-q)\left[\frac{\bar{U}-p_2}{2}\right] - c(p_2-p^*) = \frac{1+q}{2}[\bar{U}-p_2] - c(p_2-p^*)$$

respectively, where q is the probability that the rival buyer (who from this buyer's standpoint is randomly chosen from the population) visits Seller 1. So the buyer prefers to visit Seller 1 if

$$c_2 - c_1 > \frac{1+q}{2}[\bar{U}-p_2] - \left(1-\frac{q}{2}\right)[\bar{U}-p_1], \tag{7}$$

where $c_i = c(p_i - p^*)$. She prefers to visit Seller 2 if the inequality in (7) is reversed, and is indifferent if the inequality sign is replaced by an equal sign. (In particular, if $c_1 = c_2 = 0$, equality yields the equilibrium mixed–strategy under standard preferences.) So, other things equal, the larger is c_2 relative to c_1, the more likely she is to prefer Seller 1.

Now, suppose $p_1 < p_2$. Then as p^* decreases, c_2 increases relative to c_1 due to the convexity of c, so the buyer will be more likely to visit Seller 1 – who is the lower-price seller – meaning the buyer is more price–responsive. Similarly, if $p_1 > p_2$, then as p^* decreases, c_2 decreases relative to c_1, so the buyer will be more likely to visit Seller 2 (who is now the lower-price seller), so again the buyer is more price–responsive. □

Proposition 2

A seller's utility is given by

$$\Pi_i = \Phi_i \cdot p_i - c(p_i - p^*),$$

where Φ_i is the probability that at least one buyer visits the seller (which is a function of p_i and other sellers' prices, but not p^*). We assume that $\frac{\partial \Phi_i}{\partial p_i}$ and $\frac{\partial^2 \Phi_i}{\partial p_i^2}$ are negative for any rival seller price, as is true when buyers have standard preferences or weak fairness preferences, and for prices sufficiently near p^* under strong fairness preferences. We also assume that $\frac{\partial^2 \Phi_i}{\partial p_i \partial p_j} = 0$; this will be exactly true in equilibrium when buyers have standard preferences, or more generally when they have the same fairness preferences.

(i) Then a given seller's utility is maximized when

$$0 = \Phi_i + \frac{\partial \Phi_i}{\partial p_i} - \frac{\partial c}{\partial p_i},$$

and the equilibrium condition is that this holds when $p_1 = p_2 = p$, in which case $\Phi_i = 0.75$. Implicit differentiation yields

$$0 = \left[\frac{\partial \Phi_i}{\partial p_i} + p_i \frac{\partial^2 \Phi_i}{\partial p_i^2} + p_i \frac{\partial^2 \Phi_i}{\partial p_i \partial p_j} - \frac{\partial^2 c}{\partial p_i^2}\right] \cdot dp_i - \frac{\partial^2 c}{\partial p_i \partial p^*} \cdot dp^*,$$

or

$$\frac{dp}{dp^*} = \frac{-\frac{\partial^2 c}{\partial p_i \partial p^*}}{\frac{\partial^2 c}{\partial p_i^2} - \frac{\partial \Phi_i}{\partial p_i} - p\frac{\partial^2 \Phi_i}{\partial p_i^2} - p\frac{\partial^2 \Phi_i}{\partial p_i \partial p_j}} = \frac{\frac{\partial^2 c}{\partial p_i^2}}{\frac{\partial^2 c}{\partial p_i^2} - \frac{\partial \Phi_i}{\partial p_i} - p\frac{\partial^2 \Phi_i}{\partial p_i^2} - p\frac{\partial^2 \Phi_i}{\partial p_i \partial p_j}}.$$

By convexity of c, the numerator and the first term of the denominator of the last expression are positive, and by assumption, $\frac{\partial \Phi_i}{\partial p_i}$ and $\frac{\partial^2 \Phi_i}{\partial p_i^2}$ are negative and $\frac{\partial^2 \Phi_i}{\partial p_i \partial p_j}$ is zero, so that the entire denominator and thus $\frac{dp}{dp^*}$ is positive. □

(ii) Suppose now that p^* varies across sellers, with a distribution over $[0, \bar{U}]$ that is common knowledge, and that each seller's value of p^* is his private information. Then sellers' posted prices may depend on their own value of p^*.

So, a seller's expected utility is given by

$$E\Pi_i = E\Phi_i \cdot p_i - c(p_i - p^*)$$

(where the expectation is over the distribution of rival–seller prices). This is maximised when

$$0 = E\Phi_i + \frac{\partial E\Phi_i}{\partial p_i} - \frac{\partial c}{\partial p_i}.$$

Implicit differentiation yields

$$0 = \left[2\frac{\partial E\Phi_i}{\partial p_i} + p_i\frac{\partial^2 E\Phi_i}{\partial p_i^2} - \frac{\partial^2 c}{\partial p_i^2}\right] \cdot dp_i - \frac{\partial^2 c}{\partial p_i \partial p^*} \cdot dp^*,$$

or

$$\frac{dp_i}{dp^*} = \frac{-\frac{\partial^2 c}{\partial p_i \partial p^*}}{\frac{\partial^2 c}{\partial p_i^2} - 2\frac{\partial E\Phi_i}{\partial p_i} - p_i\frac{\partial^2 E\Phi_i}{\partial p^2}} = \frac{\frac{\partial^2 c}{\partial p_i^2}}{\frac{\partial^2 c}{\partial p_i^2} - 2\frac{\partial E\Phi_i}{\partial p_i} - p\frac{\partial^2 E\Phi_i}{\partial p_i^2}}.$$

By convexity of c, the numerator and the first term of the denominator of the last expression are positive, and by assumption, $\frac{\partial \Phi_i}{\partial p_i}$ and $\frac{\partial^2 \Phi_i}{\partial p_i^2}$ are negative for any rival price, so that $\frac{\partial E\Phi_i}{\partial p_i}$ and $\frac{\partial^2 E\Phi_i}{\partial p^2}$ are also negative. This means that the entire denominator and thus $\frac{dp_i}{dp^*}$ is positive. □

A.3 Solution of the 2x3–0 and 2x3–1 cases with identical risk–averse agents

We allow both sellers and buyers to have any twice–differentiable weakly risk–averse utility function such that the utility of not trading is equal to zero. The only additional assumptions we make are (a) all buyers have the same utility function and all sellers have the same (possibly different from buyers') utility function, and (b) utility functions are common knowledge amongst the agents.[22]

The 2x3–1 case

Let u be the common utility function of buyers, and v that of sellers. As noted above, the utility of not trading is zero; otherwise we express u and v as functions of p, the price at which the agent trades. Since

[22] A weaker version of (b) also works, since it is not necessary that buyers know sellers' utility functions. See Appendix A.4 for a relaxation of (a).

prices are restricted to be between zero and buyers' valuation, we have $u, v \geq 0$. Also, $v'(p) > 0 > u'(p)$ (the latter since buyer's monetary surplus is $\bar{U} - p$), and $u''(p), v''(p) \leq 0$ (from weak risk aversion). Note that the conditions on first derivatives imply that $v > 0$ when $p > 0$ and $u > 0$ when $p < \bar{U}$.

Figure 11: Buyer 1 payoffs in the 2x3–1 cell, given seller prices p_1, p_2 (symmetric game)

		Buyer 2	
		Visit Seller 1	Visit Seller 2
Buyer	Visit Seller 1	$\frac{5}{12}u_1$	$\frac{3}{4}u_1$
1	Visit Seller 2	$\frac{3}{4}u_2$	$\frac{5}{12}u_2$

Given sellers' prices, the two informed buyers play the game shown in Figure 11, where $u_1 = u(p_1)$ and $u_2 = u(p_2)$ are the utilities from buying from Seller 1 and Seller 2. (As always, the uninformed buyer, randomizes uniformly between sellers irrespective of risk attitude.) Note that each payoff is the product of the utility from buying (u_1 or u_2) and the probability of being able to buy. (For example, if Buyer 1 and Buyer 2 visit the same seller, each is equally likely to buy with probability 1/2 or 1/3, depending on the realization of the uninformed buyer's seller visit choice.) In the symmetric equilibrium of this game, each buyer visits Seller 1 with probability

$$q(p_1, p_2) = \frac{9u_1 - 5u_2}{4u_1 + 4u_2}$$

(as long as p_1 and p_2 are not too far apart), so

$$\frac{\partial q}{\partial p_1} = \frac{7u_2}{2(u_1 + u_2)^2} \cdot u'(p_1). \tag{8}$$

Given q, the probability of Seller 1 being able to sell is

$$\Phi_1(p_1, p_2) = 1 - \frac{1}{2}(1 - q)^2,$$

and Seller 1's expected utility is

$$\Pi_1(p_1, p_2) = \Phi_1(p_1, p_2) \cdot v(p_1).$$

Expected utility is maximized when

$$\begin{aligned} 0 = \frac{\partial \Pi_1}{\partial p_1} &= \Phi_1 \cdot v'(p_1) + v(p_1)\frac{\partial \Phi_1}{\partial p_1} \\ &= \Phi_1 \cdot v'(p_1) + v(p_1) \cdot (1 - q)\frac{\partial q}{\partial p_1}. \end{aligned}$$

In a symmetric equilibrium in prices, $p_1 = p_2$, so that $q = 1/2$ and thus $\Phi_1 = \Phi_2 = 7/8$. Let p be the common value of p_1 and p_2. Then p solves

$$0 = \frac{7}{8}v'(p) + v(p)\frac{7u'(p)}{16u(p)}. \tag{9}$$

Figure 12: Buyer 1 payoffs in the 2x3–0 cell, given seller prices p_1, p_2 (symmetric game)

		Buyer 2	
		Visit Seller 1	Visit Seller 2
Buyer 1	Visit Seller 1	$\frac{1}{3}u_1$	$\frac{1}{2}u_1$
	Visit Seller 2	u_2	$\frac{1}{2}u_2$

Buyer 3: Visit Seller 1

		Buyer 2	
		Visit Seller 1	Visit Seller 2
Buyer 1	Visit Seller 1	$\frac{1}{2}u_1$	u_1
	Visit Seller 2	$\frac{1}{2}u_2$	$\frac{1}{3}u_2$

Buyer 3: Visit Seller 2

The 2x3–0 case

Solving this case follows a similar process to the 2x3–1 case, with two important differences. First, all three buyers are informed, so they play a 2x2x2 game given sellers' prices (see Figure 12). In a symmetric equilibrium, each buyer visits Seller 1 with probability q that solves

$$(u_1 - u_2)q^2 - (3u_1 + u_2)q + (3u_1 - u_2) = 0,$$

so that $\frac{\partial q}{\partial p_1}$ can be computed by differentiating implicitly:

$$\frac{\partial q}{\partial p_1} = -\frac{q^2 - 3q + 3}{2q(u_1 - u_2) - 3u_1 - u_2} \cdot u'(p_1).$$

The second difference is that sellers' visit probabilities reflect the fact that all three buyers play the symmetric equilibrium strategy (rather than one of them randomizing uniformly). So, given q, we have

$$\Phi_1(p_1, p_2) = 1 - (1 - q)^3,$$

so that

$$\frac{\partial \Phi_1}{\partial p_1} = 3(1 - q)^2 \cdot \frac{\partial q}{\partial p_1}.$$

Then Seller 1's expected utility ($\Pi_1 = \Phi_1 \cdot v(p_1)$) is maximized when

$$0 = \Phi_1 \cdot v'(p_1) + v(p_1) \cdot 3(1 - q)^2 \frac{\partial q}{\partial p_1}.$$

In a symmetric equilibrium in prices, as before we have $p_1 = p_2$, $q = 1/2$ and $\Phi_1 = \Phi_2 = 7/8$, so that p (the common value of p_1 and p_2) solves

$$0 = \frac{7}{8}v'(p) + v(p)\frac{21u'(p)}{64u(p)}. \tag{10}$$

Comparison of the two cases

Let p_0 and p^* solve (10) and (9) respectively; that is, they are the subgame perfect equilibrium prices in the 2x3–0 and 2x3–1 cases given utility functions u and v. (Note that since the left–hand sides of (10) and (9) are neither zero nor infinite, it must be that u and v are strictly positive; i.e., $p_0, p^* \in (0, \bar{U})$.) Then we have

$$-2 = \frac{v(p^*)}{v'(p^*)} \cdot \frac{u'(p^*)}{u(p^*)} \quad \text{and} \quad -\frac{8}{3} = \frac{v(p_0)}{v'(p_0)} \cdot \frac{u'(p_0)}{u(p_0)},$$

so that

$$\frac{v(p^*)}{v'(p^*)} \cdot \frac{u'(p^*)}{u(p^*)} > \frac{v(p_0)}{v'(p_0)} \cdot \frac{u'(p_0)}{u(p_0)}.$$

Define

$$g(p) = \frac{u'(p)v(p)}{u(p)v'(p)}.$$

Then

$$\begin{aligned} g'(p) &= \frac{uv'(u'v' + u''v) - u'v(uv'' + u'v')}{(uv')^2} \\ &= \frac{vv'\left[uu'' - (u')^2\right] - uu'\left[vv'' - (v')^2\right]}{(uv')^2}. \end{aligned} \tag{11}$$

Since u and v are strictly positive, u'' and v'' are weakly negative, and $v' > 0 > u'$, the first term of the numerator is negative while the second is positive, making the entire numerator negative. Since the denominator is clearly positive, we therefore have $g' > 0$, so that g is increasing in p and hence $p_0 > p^*$. Thus the equilibrium price in the 2x3–0 case is higher than that in the 2x3–1 case, not only for risk–neutral agents, but also for risk–averse ones.

A.4 Solution of the 2x3–0 and 2x3–1 cases with heterogeneous risk–averse agents

Here, we extend the analysis in the previous section to allow a distribution of risk attitudes amongst buyers, with each having private knowledge of his risk attitude. To keep the analysis simple, we assume sellers are risk neutral. Uninformed buyers will continue visiting each seller with equal probability, so we focus here on informed buyers.

As in the previous section, we assume that informed buyers' utility functions are twice–differentiable and that the utility of not trading is zero. But rather than all having the same utility function, there is a one–parameter continuum of risk–averse buyer types indexed by θ, such that a buyer of type θ_1 is more risk averse than a buyer of type θ_2 if and only if $\theta_1 > \theta_2$. The utility function for a buyer of type θ will be called $u(\theta, p)$ or $u_\theta(p)$. As before, $u_\theta(p) \geq 0$, $u'_\theta(p) < 0$ and $u''_\theta(p) < 0$ for $p \in [0, \bar{U}]$ (and thus $u_\theta(p) > 0$ when $p < \bar{U}$). Suppose the population distribution of θ is given by some c.d.f. F with p.d.f. f. We assume F and f are common knowledge, but the type of any given buyer is his private information.

Suppose the sellers choose prices p_1 and p_2, where $p_1 \leq p_2$. Then visiting Seller 1 entails a smaller probability of trading but a higher surplus conditional on trading, while visiting Seller 2 entails a larger probability of trading but a lower surplus conditional on trading. This means that if buyers believe that the probability of a given buyer visiting Seller 1 is $q \geq 0.5$, and some buyer type θ is indifferent between visiting the two sellers given q, then buyers who are more risk averse (to the right of θ in the population distribution) will visit Seller 2 and those who are less risk averse (to the left of θ) will visit Seller 1, so that the fraction of buyers visiting Seller 1 will be $F(\theta)$. So, an equilibrium of the game played amongst the informed buyers will be characterized by $q^* = F(\theta^*)$, where type θ^* is indifferent between visiting the two sellers given q^*.

The 2x3–1 case

Given q, a buyer of type θ gets utility from visiting Seller 1 of $u_\theta(p_1)$ with probability $\frac{9-4q}{12}$, and gets utility from visiting Seller 2 of $u_\theta(p_2)$ with probability $\frac{5+4q}{12}$. Then, the equilibrium condition is

$$\frac{9-4q}{12} \cdot u_\theta(p_1) = \frac{5+4q}{12} \cdot u_\theta(p_2),$$

along with $q = F(\theta)$. That is,

$$[9 - 4F(\theta)] \cdot u(\theta, p_1) = [5 + 4F(\theta)] \cdot u(\theta, p_2). \tag{12}$$

To find $\frac{\partial q}{\partial p_1}$, we first note that $\frac{\partial q}{\partial p_1} = f(\theta) \cdot \frac{\partial \theta}{\partial p_1}$. Next, we use implicit differentiation on (12):

$$\begin{aligned} -4f(\theta)u(\theta, p_1)\frac{\partial \theta}{\partial p_1} \quad + \quad & [9 - 4F(\theta)]\left[u_1(\theta, p_1)\frac{\partial \theta}{\partial p_1} + u_2(\theta, p_1)\right] \\ & = 4f(\theta)u(\theta, p_2)\frac{\partial \theta}{\partial p_1} + [5 + 4F(\theta)]u_1(\theta, p_2)\frac{\partial \theta}{\partial p_1}, \end{aligned} \tag{13}$$

where u_1 and u_2 are the partial derivatives of u with respect to the first and second argument respectively.

When sellers are choosing their equilibrium prices, $p_1 = p_2$; call their common value p. Also, $q = 0.5$, which means $F(\theta) = 0.5$. So (13) simplifies to

$$-4f(\theta)u(\theta, p)\frac{\partial \theta}{\partial p_1} + 7u_2(\theta, p) = 4f(\theta)u(\theta, p)\frac{\partial \theta}{\partial p_1},$$

and solving for $\frac{\partial \theta}{\partial p_1}$ yields

$$\frac{\partial \theta}{\partial p_1} = \frac{7u'_\theta(p)}{8f(\theta)u_\theta(p)},$$

so

$$\frac{\partial q}{\partial p_1} = \frac{7u'_\theta(p)}{8u_\theta(p)}$$

when sellers are in equilibrium.

Turning to sellers, Seller 1's profit is given by $p_1 \cdot \left[1 - \frac{1}{2}(1-q)^2\right]$, which is maximized when

$$0 = \left[1 - \frac{1}{2}(1-q)^2\right] + p(1-q)\frac{\partial q}{\partial p_1}.$$

In equilibrium, $p_1 = p_2 = p$ and $q = 0.5$, so we have

$$\begin{aligned} 0 \quad &= \quad \frac{7}{8} + \frac{p}{2} \cdot \frac{\partial q}{\partial p_1} \\ &= \quad \frac{7}{8} + \frac{p}{2} \cdot \frac{7u'_\theta(p)}{8u_\theta(p)}, \end{aligned}$$

and therefore

$$-2 = p \cdot \frac{u'_\theta(p)}{u_\theta(p)},$$

which implicitly solves for the equilibrium price. (Note the similarity to the corresponding result from the previous section.)

The 2x3–0 case

Given q, a buyer of type θ gets utility $u_\theta(p_1)$ with probability $\frac{3-3q+q^2}{3}$ from visiting Seller 1, and $u_\theta(p_2)$ with probability $\frac{1+q+q^2}{3}$ from visiting Seller 2. Then the equilibrium condition is

$$[3 - 3F(\theta) + F(\theta)^2] \cdot u(\theta, p_1) = [1 + F(\theta) + F(\theta)^2] \cdot u(\theta, p_2).$$

Using implicit differentiation,

$$\begin{aligned} [-3f(\theta) + 2F(\theta)f(\theta)]u(\theta, p_1)\frac{\partial\theta}{\partial p_1} &+ [3 - 3F(\theta) + F(\theta)^2]\left[u_1(\theta, p_1)\frac{\partial\theta}{\partial p_1} + u_2(\theta, p_1)\right] \\ = [f(\theta) + 2F(\theta)f(\theta)]u(\theta, p_2)\frac{\partial\theta}{\partial p_1} &+ [1 + F(\theta) + F(\theta)^2]u_1(\theta, p_2)\frac{\partial\theta}{\partial p_1}, \end{aligned}$$

and substituting the equilibrium conditions $p_1 = p_2 = p$ and $F(\theta) = 0.5$ yields

$$-2f(\theta)u(\theta, p)\frac{\partial\theta}{\partial p_1} + \frac{7}{4}u_2(\theta, p) = 2f(\theta)u(\theta, p)\frac{\partial\theta}{\partial p_1},$$

or equivalently

$$\frac{\partial\theta}{\partial p_1} = \frac{7u'_\theta(p)}{16f(\theta)u_\theta(p)},$$

so that

$$\frac{\partial q}{\partial p_1} = \frac{7u'_\theta(p)}{16u_\theta(p)}$$

when sellers are in equilibrium.

Turning to sellers, Seller 1's profit is given by $p_1 \cdot \left[1 - (1-q)^3\right]$, which is maximized when

$$0 = \left[1 - (1-q)^3\right] + 3p(1-q)^2\frac{\partial q}{\partial p_1}.$$

In equilibrium,

$$\begin{aligned} 0 &= \frac{7}{8} + \frac{3p}{4} \cdot \frac{\partial q}{\partial p_1} \\ &= \frac{7}{8} + \frac{3p}{4} \cdot \frac{7u'_\theta(p)}{16u_\theta(p)}, \end{aligned}$$

and therefore

$$-\frac{8}{3} = p \cdot \frac{u'_\theta(p)}{u_\theta(p)}.$$

Note that the marginal buyer (the one indifferent between sellers) is the same here as in the 2x3–1 case as long as the population distribution is the same in both cases; it is the "median buyer" with type $\theta^* = F^{-1}(0.5)$.

Comparison of the two cases

To show that the equilibrium price must be higher in 2x3–0 than in 2x3–1, we follow a process very much like that at the end of Appendix A.3 for the homogeneous case. Define $h(p) = p \cdot \frac{u'_\theta(p)}{u_\theta(p)}$; then as before, it suffices to show that $h'(p) < 0$. But

$$\begin{aligned} h'(p) &= \frac{u_\theta(p)[p \cdot u''_\theta(p) + u'_\theta(p)] - p[u'_\theta(p)]^2}{[u_\theta(p)]^2} \\ &= \frac{pu_\theta(p)u''_\theta(p) + u_\theta(p)u'_\theta(p) - p[u'_\theta(p)]^2}{[u_\theta(p)]^2}. \end{aligned} \tag{14}$$

The denominator of (14) is positive, while the three terms of the numerator are negative, negative and positive respectively, so that the entire fraction is negative. Thus even under heterogeneous and privately–known risk attitudes, the predicted treatment effect remains the same.

A.5 Effect of price responsiveness on equilibrium prices

As in Section 4.4, suppose that buyers are β times as price–responsive as they are predicted to be under the assumption of standard preferences (i.e., suppose that price responsiveness is given by (3) in that section), while sellers have standard preferences. In this section, we compute the implied Nash equilibrium in seller prices, as a function of β, for the 2x2–0 case. The 2x3–0 and 2x3–1 cases follow a nearly identical process (and indeed can be obtained from Appendix A.3, substituting in the CRRA utility function $u(x) = x^\beta$ but with $v(x) = x$), while the 2x2–1 case is analogous to what was done in Appendix A.1.[23]

A seller's expected profit is equal to her price multiplied by the probability that she is visited by at least one buyer. For Seller 1 in the 2x2–0 case, this is $\pi_1 = p_1 \cdot \Phi$, where

$$\Phi = 1 - (1 - \hat{q})^2, \tag{15}$$

$\hat{q}$ is given by (3), and $q(p_1, p_2)$ is given by (2) in Section 4.3. Then

$$\begin{aligned} \frac{\partial \pi_1}{\partial p_1} &= 1 - (1 - \hat{q})^2 + 2p_1(1 - \hat{q})\frac{\partial \hat{q}}{\partial p_1} \\ &= 1 - (1 - \hat{q})^2 + 2p_1(1 - \hat{q})\beta\frac{\partial q}{\partial p_1} \\ &= 1 - (1 - \hat{q})^2 + 2p_1(1 - \hat{q})\beta\left[-\frac{60 - 3p_2}{(40 - p_1 - p_2)^2}\right]. \end{aligned} \tag{16}$$

In a symmetric Nash equilibrium, π_1 is maximized when $p_1 = p_2$, so that q and hence $\hat{q}$ are equal to one–half. So (16) simplifies to $0 = \frac{3}{4} - \frac{3\beta p}{4(20-p)}$ (where p is the common value of p_1 and p_2), whence $p = \frac{20}{1+\beta}$.

[23] The associated equilibrium prices are $\frac{160}{8+3\beta}$ in the 2x3–0 case and $\frac{40}{2+\beta}$ in the 2x3–1 case, while in the 2x2–1 case, sellers play continuous mixed strategies with expectation $\frac{30-10\beta}{\beta} \cdot ln\left(\frac{3+\beta}{3-\beta}\right)$, based on the density function $f(p) = \frac{30-10\beta}{\beta p^2}$ for $p \in \left[20\left(\frac{3-\beta}{3+\beta}\right), 20\right]$.

A.6 Sample instructions from original experiment

The instructions below are from the 2x3 market, with no uninformed buyers in rounds 1-20 and one uninformed buyer in rounds 21-40. Parts 1, 2 and 3 were given to subjects before round 1, between rounds 20 and 21, and after round 40 respectively. The instructions from the other treatments were very similar, and available from the corresponding author upon request.

Instructions (Part 1)

You are about to participate in a decision making experiment. Please read these instructions carefully, as the money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment, and that you put away your mobile phones and other devices at this time.

This experiment has three parts. These instructions are for Part 1; you will receive new instructions after this part has finished. Part 1 is made up of 20 rounds, each consisting of a simple computerised market game. Before the first round, you are assigned a role: buyer or seller. ***You will remain in the same role throughout the experiment.***

In each round, the participants in this session are divided into "markets": groups of five containing a total of ***three buyers*** and ***two sellers***. ***The other people in your market will be randomly assigned in each round.*** You will not be told the identity of the people in your market, nor will they be told yours – even after the session ends.

The market game: In each round, a seller can produce ***one*** unit of a hypothetical good, at a cost of ***$0***. A buyer can buy up to ***one*** unit of the good, which is ***resold to the experimenter*** at the end of the round for ***$20***. It is not possible to buy or sell more than one unit in a round. Sellers begin by choosing their prices, which are entered as ***multiples of 0.05, between 0 and 20*** inclusive (without the dollar sign).

After all sellers have chosen prices, ***all buyers can observe the prices*** of each of the sellers in their market, then each chooses which seller to visit. If ***only one buyer*** visits a particular seller, then that buyer pays the seller's price for the seller's item. If ***more than one buyer*** visits the ***same*** seller, then since the seller only has one unit, ***one of the buyers is randomly selected*** by the computer to purchase it at that seller's price, and the other buyers are ***unable to buy***. Any seller visited by no buyers is ***unable to sell***.

Profits: Your profit for the round depends on the round's result.
- If you are a ***seller*** and you are ***able to sell***, your profit is the price you chose.
- If you are a ***seller*** and you are ***unable to sell***, your profit is zero.
- If you are a ***buyer*** and you are ***able to buy***, your profit is $20.00 minus the price you paid.
- If you are a ***buyer*** and you are ***unable to buy***, your profit is zero.

Sequence of play in a round:
(1) The computer randomly forms markets made up of ***three buyers*** and ***two sellers*** each.
(2) Sellers choose their prices.
(3) All buyers observe the sellers' prices, then each buyer chooses which seller to visit.
(4) The round ends. If you are a seller, you are informed of: each seller's price, how many buyers visited

you, quantity sold and profit for the round. If you are a buyer, you are informed of: each seller's price, how many buyers visited the same seller as you, your quantity bought and profit for the round. After this, you go on to the next round.

Payments: Your payment depends on the results of the experiment. At the end of the experiment, ***two*** rounds from Part 1 will be chosen randomly for each participant. ***You will be paid the total of your profits from those two rounds***, plus whatever you earn in the other parts of the experiment. Payments are made privately and in cash at the end of the session.

Instructions (Part 2)

The procedure in this part is nearly the same as before. You will play a similar market game for 20 additional rounds. Your role (buyer or seller) will remain the same as before, and the participants in your market – a total of 3 buyers and 2 sellers – will still be randomly re-assigned in each round.

The difference from the first part of the experiment is that now, in each round ***only two of the three buyers will be able to observe the sellers' prices*** before choosing which seller to visit. The remaining buyer will still choose a seller to visit, but ***without having seen any prices***.

As before, if only one buyer visits a particular seller, then that buyer buys the seller's item, and if more than one buyer visits the same seller, then one of them is randomly selected by the computer to buy. Profits are determined in the same way as in Part 1.

Payments: At the end of the experiment, ***two*** rounds from Part 2 will be chosen randomly for each participant. ***You will be paid the total of your profits from those two rounds***. Your payment from this part will be added to what you got from Part 1, and additional earnings are possible in Part 3.

Instructions (Part 3)

This part begins with a Gamble Selection Task. You will be shown five gambles, and will be asked to choose the one you prefer. Each gamble has two possible outcomes, with equal (50%) chance of occurring. Your earnings from this task will depend on which gamble you choose, and which outcome occurs.

The gambles are as follows:

Gamble	Random numbers 1-50 (50% chance)	Random numbers 51-100 (50% chance)
1	You earn $4	You earn $4
2	You earn $6	You earn $3
3	You earn $8	You earn $2
4	You earn $10	You earn $1
5	You earn $12	You earn $0

After you have chosen one of these gambles, the computer will randomly draw a whole number between 1 and 100 (inclusive). If the random number is 50 or less, your earnings from this task are as shown in

the middle column of the table. If the random number is 51 or more, your earnings from this task are as shown in the right column. The random number drawn for you may be different from the ones drawn for other participants.

Once you have chosen a gamble, you will be shown another screen containing a questionnaire. You will receive $5 for answering all of the questions if you were a seller in Parts 1 and 2, or $15 if you were a buyer. Once everyone has finished the questionnaire, you will be shown the results of all three parts of the experiment.

A.7 Sample screen-shots

Seller decision screen:

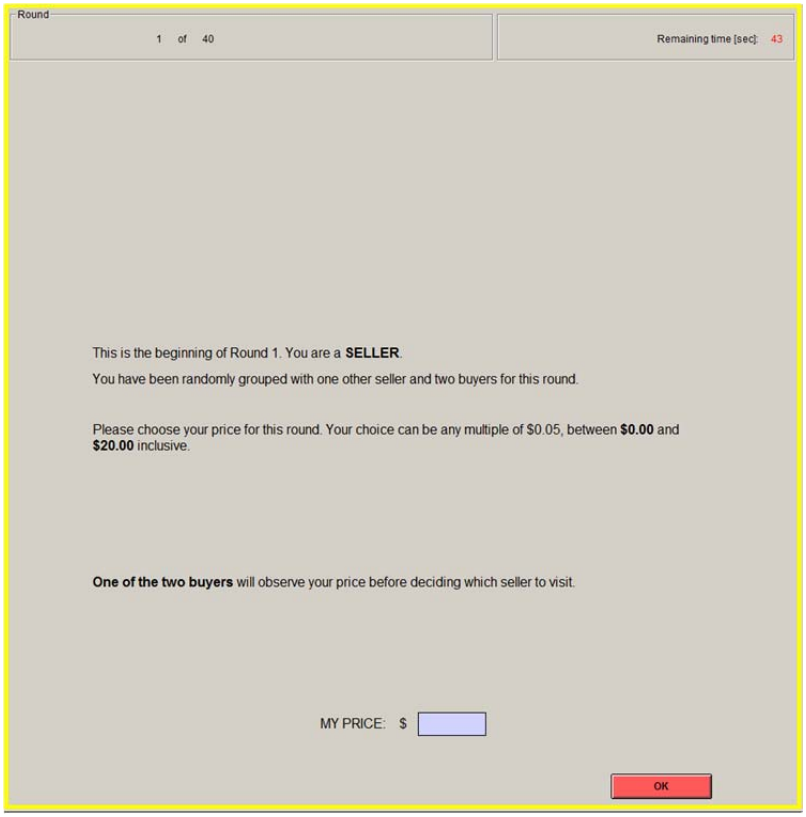

Informed buyer decision screen:

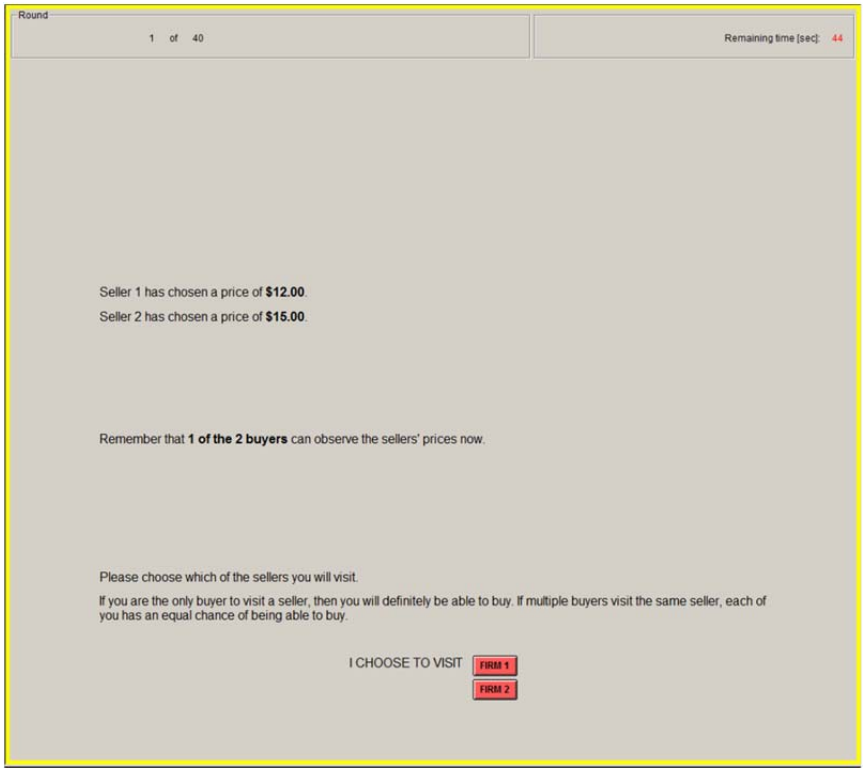

Uninformed buyer decision screen:

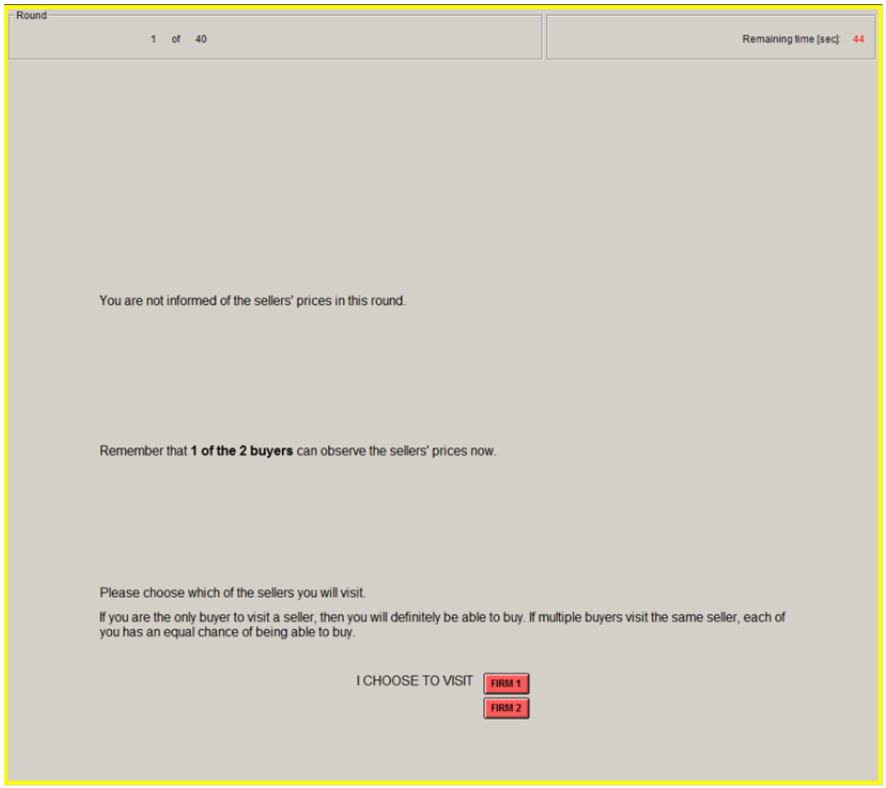

Lottery choice screen:

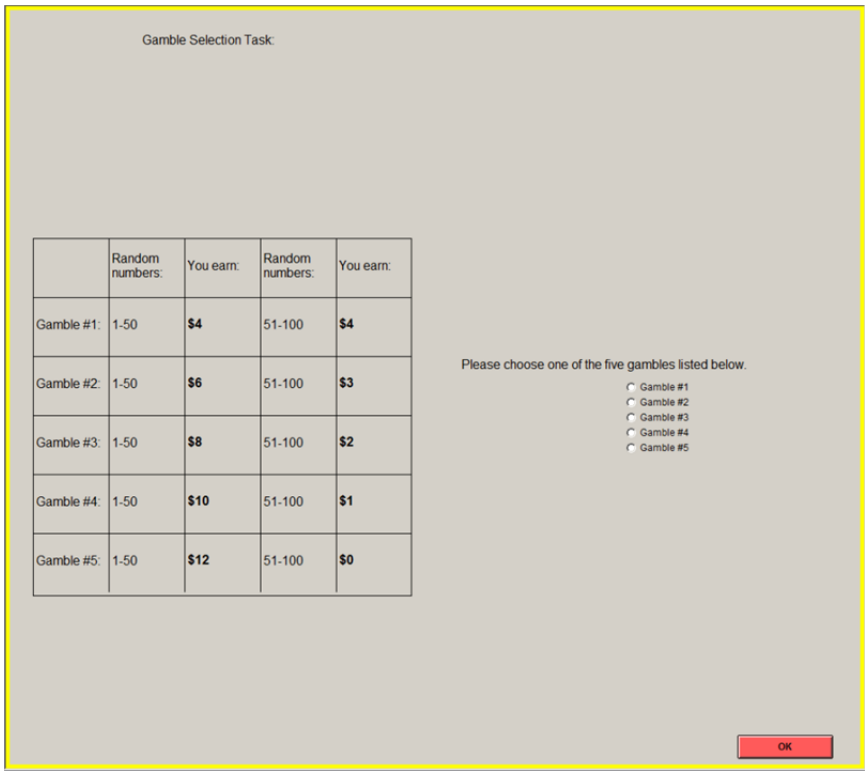

	Random numbers:	You earn:	Random numbers:	You earn:
Gamble #1:	1-50	**$4**	51-100	**$4**
Gamble #2:	1-50	**$6**	51-100	**$3**
Gamble #3:	1-50	**$8**	51-100	**$2**
Gamble #4:	1-50	**$10**	51-100	**$1**
Gamble #5:	1-50	**$12**	51-100	**$0**

Questionnaire screen:

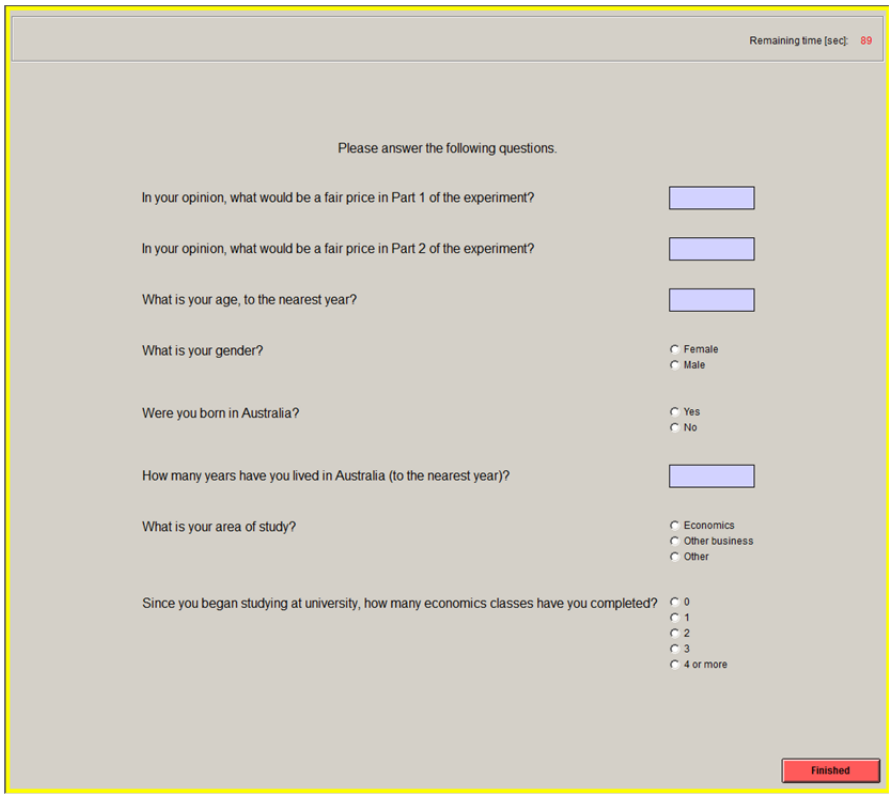

A.8 Sample instructions from follow-up experiment

The instructions below are from the 2x3 market with one uninformed buyer in all rounds (1-50). Parts 1 and 2 were given to subjects before round 1 and after round 50 respectively. The instructions from the other treatment were very similar, and available from the corresponding author upon request.

Instructions (Part 1)

You are about to participate in a decision making experiment. Please read these instructions carefully, as the money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment, and that you put away your mobile phones and other devices at this time.

This experiment has two parts. These instructions are for Part 1; you will receive new instructions after this part has finished. Part 1 is made up of 50 rounds, each consisting of a simple computerised market game. Before the first round, you are assigned a role: buyer or seller. ***You will remain in the same role throughout the experiment.***

In each round, the participants in this session are divided into "markets": groups of five containing a total of ***three buyers*** and ***two sellers***. ***The other people in your market will be randomly assigned in each round.*** You will not be told the identity of the people in your market, nor will they be told yours – even after the session ends.

The market game: In each round, a seller can produce ***one*** unit of a hypothetical good, at a cost of ***$0***. A buyer can buy up to ***one*** unit of the good, which is ***resold to the experimenter*** at the end of the round for ***$20***. It is not possible to buy or sell more than one unit in a round. Sellers begin by choosing their prices, which are entered as ***multiples of 0.05, between 0 and 20*** inclusive (without the dollar sign).

After all sellers have chosen prices, ***two of the three buyers can observe the prices*** of each of the sellers in their market, while the remaining buyer does not. Then each buyer chooses which seller to visit. If ***only one buyer*** visits a particular seller, then that buyer pays the seller's price for the seller's item. If ***more than one buyer*** visits the ***same*** seller, then since the seller only has one unit, ***one of the buyers is randomly selected*** by the computer to purchase it at that seller's price, and the other buyers are ***unable to buy***. Any seller visited by no buyers is ***unable to sell***.

Profits: Your profit for the round depends on the round's result.
- If you are a ***seller*** and you are ***able to sell***, your profit is the price you chose.
- If you are a ***seller*** and you are ***unable to sell***, your profit is zero.
- If you are a ***buyer*** and you are ***able to buy***, your profit is $20.00 minus the price you paid.
- If you are a ***buyer*** and you are ***unable to buy***, your profit is zero.

Sequence of play in a round:
(1) The computer randomly forms markets made up of ***three buyers*** and ***two sellers*** each.
(2) Sellers choose their prices.
(3) Two of the three buyers observe the sellers' prices, then each buyer chooses which seller to visit.
(4) The round ends. If you are a seller, you are informed of: each seller's price, how many buyers visited you, quantity sold and profit for the round. If you are a buyer, you are informed of: each seller's price, how many buyers visited the same seller as you, your quantity bought and profit for the round.

After this, you go on to the next round.

Payments: Your payment depends on the results of the experiment. At the end of the session, ***five*** rounds from Part 1 will be chosen randomly for each participant. ***You will be paid the total of your profits from those five rounds***, plus whatever you earn in Part 2 of the experiment. Payments are made privately and in cash at the end of the session.

Instructions (Part 2)

This part begins with a Gamble Selection Task. You will be shown five gambles, and will be asked to choose the one you prefer. Each gamble has two possible outcomes, with equal (50%) chance of occurring. Your earnings from this task will depend on which gamble you choose, and which outcome occurs.

The gambles are as follows:

Gamble	Random numbers 1-50 (50% chance)	Random numbers 51-100 (50% chance)
1	You earn $4	You earn $4
2	You earn $6	You earn $3
3	You earn $8	You earn $2
4	You earn $10	You earn $1
5	You earn $12	You earn $0

After you have chosen one of these gambles, the computer will randomly draw a whole number between 1 and 100 (inclusive). If the random number is 50 or less, your earnings from this task are as shown in the middle column of the table. If the random number is 51 or more, your earnings from this task are as shown in the right column. The random number drawn for you may be different from the ones drawn for other participants.

Once you have chosen a gamble, you will be shown another screen containing a questionnaire. You will receive $5 for answering all of the questions if you were a seller in Part 1, or $15 if you were a buyer. Once everyone has finished the questionnaire, you will be shown the results of both parts of the experiment.