11-27-2019

Price Signaling and Bargain Hunting in Markets with Partially Informed Populations

Mark Schneider
Chapman University, maschneider4@cba.ua.edu

Daniel Graydon Stephenson
Chapman University, stephensod@vcu.edu

Price Signaling and Bargains in Markets with Partially Informed Populations*

Mark Schneider
maschneider4@cba.ua.edu

Daniel Graydon Stephenson
stephensod@vcu.edu

November 27, 2019

Abstract

Classical studies of asymmetric information focus on situations where only one side of a market is informed. This study experimentally investigates a more general case where some sellers are informed and some buyers are informed. We establish the existence of semi-separating perfect Bayesian equilibria where prices serve as informative signals of quality to uninformed buyers, while informed buyers can often leverage their informational advantage by purchasing high quality items from uninformed sellers at bargain prices. These models provide a rational foundation for the co-existence of bargains, price signaling, and Pareto efficiency in markets with asymmetric information. We test these theoretical predictions in a controlled laboratory experiment where agents repeatedly participate in markets with asymmetric information. We observe long run behavior consistent with equilibrium predictions of price signaling, bargains, and partial-pooling behavior.

*We thank George Akerlof, Manel Baucells, the seminar participants at the Darden School of Business, and the conference participants at the 2019 ESA North America meetings for their insightful comments on this paper.

1 Introduction

Economic models frequently assume that consumers and producers have perfect information about the quality of items in the marketplace. In contrast to this assumption, more recent work considers the case of asymmetric information where some agents have more information than others. Marketplaces such as Amazon, AbeBooks, eBay, flea markets, and trade shows are often characterized by significant uncertainty regarding the quality of the goods available for sale.

Akerlof (1970) considers an information structure in which all sellers are informed about the quality of the items they offer but all buyers are uninformed about the quality of the items they are offered. Here buyers can only observe the average quality of items in the marketplace. Akerlof's classic analysis derived the implication that markets will unravel such that higher quality items will disappear from the marketplace leaving only 'lemons' for sale.

Casual observation suggests that some consumers can find bargains by purchasing high quality items at unusually low prices. One hears of people who find valuable paintings at a flea market or who find rare first editions at a library book sale. Some consumers go to these venues primarily to look for such 'good deals.' These types of bargains are difficult to rationalize in Akerlof's setting since only lemons are available for sale in equilibrium.

Like Akerlof (1970), Bagwell and Riordan (1991), and Janssen and Roy (2010), we adopt a setting with binary quality. In our setting, some sellers are uninformed such that they cannot distinguish high quality items from low quality items. In such markets, experienced collectors and connoisseurs may have detailed knowledge of the items they collect, allowing them to take advantage of underpriced items sold by uninformed sellers.

The possibility that some sellers are uninformed is especially plausible in markets where evaluating goods requires some expertise or where there is a cost to becoming informed. Valuation of artwork, collectibles, or antiques often requires expertise. Researchers at eBay Research Labs (Hu and Bolivar (2008)) found that the average consumer surplus ratio in eBay auctions for items in the collectibles category is approximately 40% and that for certain subcategories such as Pre-1940 photographic image collectibles, the median consumer surplus ratio was over 50%. If sellers were informed about the value of their items, they could have extracted a higher profit from setting a higher reserve price or opening bid. In contrast, iPhones, which are easier to value, were found to yield a consumer surplus ratio of 1.59%.

Sellers with large inventories may also be uninformed about the value of individual items and may find it prohibitively costly to examine and value each item. For instance, a seller with an inventory of thousands of used books may believe that some of the books could be valuable first editions, but consider such books too rare to justify inspecting the entire inventory. Even in markets where sellers are expected to be informed such as real estate markets, recent studies find evidence that many sellers are uninformed about the energy efficiency of their homes (Cassidy (2019); Myers et al. (2019)).

In practice, uninformed consumers often try to infer quality directly from the price of a product. For example, wine consumers often assume that high quality products will be offered at higher prices. The possibility that price might serve as an informative signal of quality does not typically arise in neoclassical general equilibrium theory where item qualities are commonly known, nor does it arise in Akerlof's model where only low-quality items are available for sale in equilibrium.

Previous work has considered special cases where some buyers are informed or some sellers are informed. Prior work by Salop and Stiglitz (1977) studied markets with bargains and ripoffs but made the strong assumption that all items in the market have the same quality. Their model does not address the problem of adverse selection that is central to Akerlof's analysis.

Subsequent work by Chan and Leland (1982), Wolinsky (1983), and Cooper and Ross (1984) showed that price signaling can arise in equilibrium if some buyers are informed and firms can choose the quality level they produce. Bagwell and Riordan (1991) and Janssen and Roy (2010) demonstrate that price can also signal quality when a firm's level of quality is given and all sellers are informed. Kessler (2001) considers a perfectly competitive market where some sellers are uninformed, all buyers are uninformed, and all agents act as price-takers. Dari-Mattiacci et al. (2011) analyze markets where all buyers are informed and all sellers are uninformed. They find 'inverse adverse selection' in which the market disappears from the bottom rather than from the top.

Nearly five decades later, the strategic analysis of markets with asymmetric information remains incomplete. The previous literature has surprisingly not considered the case where there are both informed and uninformed agents on each side of the market. Here we study this more general case in a non-cooperative signaling game in which sellers strategically choose prices, and buyers endogenously form beliefs. In this paper, we prove the existence of perfect Bayesian Nash equilibria where uninformed sellers either pool with informed high quality sellers or informed low quality sellers.

To test these theoretical predictions, we conduct laboratory experiments where heterogeneous buyers and sellers repeatedly participate in a market for goods with heterogeneous quality. Consistent with the equilibrium predictions, we find evidence of both bargains and price-signaling behavior. We also find evidence consistent with the predicted pooling behavior of uninformed sellers. Observed behavior is well-explained by a noisy best-response model that closely approximates the equilibrium predictions.

We characterize two types of perfect Bayesian Nash equilibria: low price equilibria, where uninformed sellers charge low prices, leading to Pareto efficient full trade; and high-price equilibria, where uninformed sellers charge high prices. If enough items are low quality, uninformed sellers are better off charging low prices. In this case, high prices are informative signals of high quality since only informed high quality sellers post them. Thus, uninformed buyers will be willing to purchase items with high prices. This case forms the basis for our low price equilibrium.

Conversely, if enough buyers are informed, uninformed sellers should primarily respond to informed buyers. If enough items are high quality, uninformed sellers are better off charging high prices. In this case, high prices are not reliable signals of high quality. Hence uninformed buyers will not be willing to pay high prices. This case forms the basis for our high price equilibrium.

In the low price equilibrium, informed buyers can purchase high quality items from uninformed sellers at bargain prices. These bargains are rare because most items are low quality. This is why uninformed sellers are better off setting a low price and selling to everyone rather than setting a high price and only selling to informed buyers if their item is high quality. So, in equilibrium, bargains can only occur when they are sufficiently rare. This mirrors a feature of bargains in real markets where informed consumers search through large numbers of low quality offers in hopes of finding a high quality item at a low price.

Our experiment consists of two treatments. Adaptive dynamics predict convergence to low price equilibria in our first treatment. Conversely, adaptive dynamics predict convergence to high price equilibria in our second treatment. The observed behavior supports these qualitative predictions: In the first treatment, uninformed sellers largely pooled with low quality informed sellers. In the second treatment uninformed sellers largely pooled with high quality informed sellers.

While subject behavior was largely consistent with these pooling predictions, it systematically deviated from equilibrium predictions in terms of the transaction rate and the bargain rate. Equilibrium predicts full trade in our first treatment but low levels of trade in our second treatment. While the first treatment exhibited significantly higher empirical transaction rates, it remained significantly below equilibrium predictions. Equilibrium predicts that uninformed sellers will always offer bargains in our first treatment but will never offer bargains in our second treatments. While the first treatment exhibited significantly higher empirical bargain rates, it remained significantly below equilibrium predictions.

The remaining sections are organized as follows. Section 2 formally introduces the signaling game and derives equilibrium predictions. Section 3 describes the experimental design and procedures. Section 4 identifies the hypotheses to be tested. Section 5 presents the results and Section 6 concludes. All proofs are included in Appendix A.

2 Theory

Consider the following interaction between a buyer and a seller. The seller possesses an item which she values at $q \in Q = \{\underline{q}, \bar{q}\}$ such that $\underline{q} < \bar{q} \in \mathbb{R}_+$. With probability θ, the item is low quality ($q = \underline{q}$). With probability $1 - \theta$, the item is high quality ($q = \bar{q}$). With probability $\lambda \in (0, 1)$, the seller is uninformed about the quality of the item. With probability $1 - \lambda$, the seller is informed about the quality of the item. Let I_s denote the seller's information level such that $I_s = 1$ if the seller is informed and $I_s = 0$ if the seller is uninformed.

The buyer values the seller's item at kq such that $k > 1$. With probability $\gamma \in (0, 1)$, the buyer is uninformed about the quality of the item. With probability $1 - \gamma$ the buyer is informed about the quality of the item. Let I_b denote the buyer's information level such that $I_b = 1$ if the buyer is informed and $I_b = 0$ if the buyer is uninformed. The seller chooses a posted price $p \in \mathbb{R}_+$ for the item. After observing the posted price, the buyer decides whether to purchase the item. Let $B = 1$ if the buyer decides to purchase the item and $B = 0$ if the buyer decides not to purchase the item.

2.1 Strategies

Let Ω denote the state space such that

$$\Omega = \left\{(q, I_b, I_s) : q \in \left\{\bar{q}, \underline{q}\right\}, I_b, I_s \in \{0, 1\}\right\} \tag{1}$$

We say that an uninformed seller has type U_s. We say that an informed seller with a high quality item has type H_s. We say that an informed seller with a low quality item has type L_s. Let $T_s = \{H_s, U_s, L_s\}$ denote the seller's type set. The seller's type $t_s \in T_s$ is given by

$\tau_s : \Omega \to T_s$ such that

$$\tau_s(\omega) = \begin{cases} U_s & \text{if } I_s = 0 \\ H_s & \text{if } I_s = 1 \text{ and } q = \bar{q} \\ L_s & \text{if } I_s = 1 \text{ and } q = \underline{q} \end{cases} \tag{2}$$

The seller's strategy is given by $\rho : T_s \to \mathbb{R}_+$ such that $p = \rho(t_s)$. That is, the seller's strategy specifies her posted price as a function of her type.

We say that an uninformed buyer has type U_b. We say that an informed buyer considering a high quality item has type H_b. We say that an informed buyer considering a low quality item has type L_b. Let $T_b = \{H_b, U_b, L_b\}$ denote the buyer's type set. The buyer's type $t_b \in T_b$ is given by $\tau_b : \Omega \to T_b$ such that

$$\tau_b(\omega) = \begin{cases} U_b & \text{if } I_b = 0 \\ H_b & \text{if } I_b = 1 \text{ and } q = \bar{q} \\ L_b & \text{if } I_b = 1 \text{ and } q = \underline{q} \end{cases} \tag{3}$$

The buyer's strategy is given by $\beta : T_b \times \mathbb{R}_+ \to \{0, 1\}$ such that

$$B = \beta(t_b, p) \tag{4}$$

That is, the buyer's strategy specifies whether or not she will buy an item as a function of her type and the posted price of the item. In some cases, the buyer's strategy may take the form of a reservation price $R : T_b \to \mathbb{R}_+$ such that

$$\beta(t_b, p) = \begin{cases} 1 & p \leq R(\tau_b) \\ 0 & p > R(\tau_b) \end{cases} \tag{5}$$

2.2 Expected Payoffs

If the buyer decides to purchase the item then she pays the posted price to the seller and receives the item from the seller. Let π_b denote the buyer's payoff and π_s denote the seller's payoff such that

$$\pi_b = \beta(\tau_b, p)(kq - p) \tag{6}$$

$$\pi_s = \beta(\tau_b, p)(p - q) \tag{7}$$

Since $k > 1$, there are always gains from trade. An informed seller with a high quality item knows that an informed buyer must also observe high quality. Conversely an informed buyer considering a high quality item knows that an informed seller must also observe high quality. Let T denote the feasible type space such that

$$T = T_b \times T_s \setminus \{(H_b, L_s), (L_b, H_s)\}$$

Here, T denotes all possible type profiles such that informed agents agree about the quality of the item. The conditional distribution of the buyer's type t_b given the seller's type t_s is

$\mathbb{P}(t_b \mid t_s)$	$t_b = H_b$	$t_b = U_b$	$t_b = L_b$
$t_s = H_s$	$1-\gamma$	γ	0
$t_s = U_s$	$(1-\theta)(1-\gamma)$	γ	$\theta(1-\gamma)$
$t_s = L_s$	0	γ	$1-\gamma$

Let q_U denote the unconditional expected quality ($q_U = \theta \underline{q} + (1-\theta)\bar{q}$). In contrast, the conditional expected quality of the item given the type profile $t = (t_b, t_s) \in T$ is

$$\mathbb{E}\{q \mid t_b, t_s\} = \begin{cases} \bar{q} & \text{if } t_b = H_b \text{ or } t_s = H_s \\ q_U & \text{if } t_b = U_b \text{ and } t_s = U_s \\ \underline{q} & \text{if } t_b = L_b \text{ or } t_s = L_s \end{cases} \tag{8}$$

Hence, the seller's conditional expected payoff given her type t_s and her posted price p is

$$\mathbb{E}\{\pi_s \mid t_s, p\} = \sum_{t_b \in T_b} \mathbb{P}(t_b \mid t_s)\, \beta(t_b, p)\,(p - \mathbb{E}\{q \mid t_b, t_s\}) \tag{9}$$

Let $\mu(q \mid p, t_b)$ denote the buyer's posterior belief regarding the quality of the item conditional on her type t_b and the posted price p. The conditional expected quality given the buyer's type and the posted price of the item is

$$\mathbb{E}_\mu\{q \mid t_b, p\} = \mu(\bar{q} \mid p, t_b)\,\bar{q} + \mu\left(\underline{q} \mid p, t_b\right)\underline{q} \tag{10}$$

The buyer's expected payoff conditional on her type t_b, the posted price p, and her purchasing decision $B \in \{0, 1\}$ is given by

$$\mathbb{E}_\mu\{\pi_b \mid t_b, p, B\} = (k\mathbb{E}_\mu\{q \mid t_b, p\} - p)\, B \tag{11}$$

2.3 Perfect Bayesian Equilibria

The conditional probability of observing a posted price p given the seller's type t_s is given by

$$\mathbb{P}\left(p|\tau_s\right) = \begin{cases} 1 & \text{if } p = \rho\left(\tau_s\right) \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

Hence the unconditional probability of observing a posted price p is given by

$$\mathbb{P}\left(p\right) = \lambda\mathbb{P}\left(p|U_s\right) + \theta\left(1-\lambda\right)\mathbb{P}\left(p|L_s\right) + \left(1-\theta\right)\left(1-\lambda\right)\mathbb{P}\left(p|H_s\right) \tag{13}$$

The conditional probability of observing a posted price p given the item's quality q is

$$\mathbb{P}\left(p|q\right) \quad = \quad \begin{cases} \lambda\mathbb{P}\left(p|U_s\right) + \left(1-\lambda\right)\mathbb{P}\left(p|L_s\right) & \text{if } q = \underline{q} \\ \lambda\mathbb{P}\left(p|U_s\right) + \left(1-\lambda\right)\mathbb{P}\left(p|H_s\right) & \text{if } q = \bar{q} \end{cases} \tag{14}$$

Now if $\mathbb{P}\left(p\right) > 0$ then Bayes' rule implies that

$$\mathbb{P}\left(\underline{q}|p\right) = \frac{\mathbb{P}\left(\underline{q}\right)\mathbb{P}\left(p|\underline{q}\right)}{\mathbb{P}\left(p\right)} = \frac{\theta\mathbb{P}\left(p|\underline{q}\right)}{\mathbb{P}\left(p\right)}$$

Hence the buyer's belief μ is consistent along the path of play if

$$\mu\left(\underline{q}|p,t_b\right) = \begin{cases} 1 & \text{if } t_b = L_b \\ 0 & \text{if } t_b = H_b \\ \mathbb{P}\left(\underline{q}|p\right) & \text{if } t_b = U_b \text{ and } \mathbb{P}\left(p\right) > 0 \end{cases} \tag{15}$$

Equation (15) simply states that the buyer believes the item is low quality if she knows that it is low quality, she believes that it is high quality if she knows it is high quality, and she uses Bayes' rule when possible. A strategy profile (ρ, β) and a consistent belief μ form a perfect Bayesian equilibrium if

$$\rho\left(t_s\right) \in \underset{p\in\mathbb{R}_+}{\operatorname{argmax}}\ \mathbb{E}\left\{\pi_s|t_s,p\right\} \quad \text{for } t_s \in T_s$$

$$\beta\left(t_b,p\right) \in \underset{B\in\{0,1\}}{\operatorname{argmax}}\ \mathbb{E}_\mu\left\{\pi_b|t_b,p,B\right\} \quad \text{for } \left(t_b,p\right) \in T_s \times \mathbb{R}_+$$

2.3.1 No Trade Equilibria

If the seller's price is always greater than the buyer's value for a high quality item then the buyer might as well refuse to even consider the item. On the other hand, if the buyer refuses to even consider the item, then the seller might as well post a very high price. Formally, if $\rho(t_s) \geq k\bar{q}$ for all $t_s \in T_s$ then $k\mathbb{E}_\mu \{q|t_b, \rho(t_s)\} \leq \rho(t_s)$ for all $(t_b, t_s) \in T$. Conversely, if $\beta(t_b, p) = 0$ for all $p \in \mathbb{R}_+$ then $\mathbb{E}\{\pi_s|t_s, p\} = 0$ for all $(t_s, p) \in T_s \times \mathbb{R}_+$. Hence there exists a perfect Bayesian equilibrium under which a transaction never occurs as $\beta(t_b, \rho(t_s)) = 0$ for all $(t_b, t_s) \in T$.

2.3.2 Trade Equilibria

A perfect Bayesian equilibrium (ρ, β) is said to be a trade equilibrium if

$$\beta(\rho(L_s), L_b) = \beta(\rho(H_s), H_b) = 1 \tag{16}$$

In a trade equilibrium, (16) indicates that a transaction always takes place if the buyer values the item more than the seller and the quality of the item is known to both parties. Proposition 1 states that the uninformed seller pools with either the informed high quality seller or the informed low quality seller in every trade equilibrium.

Proposition 1. *If (ρ, β) is a trade equilibrium then $\rho(U_s) \in \{\rho(L_s), \rho(H_s)\}$.*

2.3.3 Strategy Profiles

We say that a strategy profile (ρ, β) is of type 1 if

$$\beta(t_b, p) = \begin{cases} 1 & p \leq R(t_b) \\ 0 & p > R(t_b) \end{cases} \tag{17}$$

$$\underline{p} = \rho(L_s) = \rho(U_s) < \rho(H_s) = \bar{p} \tag{18}$$

$$\underline{p} = R(L_b) < R(U_b) = R(H_b) = \bar{p} \tag{19}$$

Under a type 1 strategy profile, the seller posts a low price $\underline{p}$ if she is uninformed or she is informed about a low quality item. She posts a high price $\bar{p}$ if she is informed about a high quality item. The buyer is willing to pay either price if she is uninformed or she is informed about a high quality item. The buyer is willing to pay the low price $\underline{p}$ but not the high price $\bar{p}$ if she is informed about a low quality item.

We say that a strategy profile (ρ, β) is of type 2 if

$$\beta(t_b, p) = \begin{cases} 1 & p \leq R(\tau_b) \\ 0 & p > R(\tau_b) \end{cases} \tag{20}$$

$$\underline{p} = \rho(L_s) < \rho(U_s) = \rho(H_s) = \bar{p} \tag{21}$$

$$\underline{p} = R(L_b) = R(U_b) < R(H_b) = \bar{p} \tag{22}$$

Under a type 2 strategy profile, the seller posts a high price $\bar{p}$ if she is uninformed or she is informed about a high quality item. She posts a low price $\underline{p}$ if she is informed about a low quality item. The buyer is willing to pay either price if she is informed about a high quality item. The buyer is willing to pay the low price $\underline{p}$ but not the high price $\bar{p}$ if she is informed about a low quality item or if she is uninformed.

2.3.4 Low Price Equilibria

Proposition 2 states that a broad class of environments support low price equilibria with type 1 strategy profiles. Equation (23) and equation (24) are the individual rationality constraints that an informed buyer and an informed seller each guarantee themselves a non-negative payoff. They also imply the possibility of gains from trade between the buyer and the seller for both low quality items and high quality items. Equation (25) is the incentive compatibility constraint that an informed seller with a low quality item is better off selling at the low price with certainty than selling at the high price only if the buyer is uninformed. Equation (26) is the incentive compatibility constraint that an uninformed seller is better off selling at the low price with certainty than selling at the high price only if either the quality is high and the buyer is informed or if the buyer is uninformed.

Proposition 2. *If (ρ, β) is a type 1 strategy profile such that*

$$\bar{q} \leq \bar{p} \leq k\bar{q} \tag{23}$$

$$\underline{q} \leq \underline{p} \leq k\underline{q} \tag{24}$$

$$\gamma\left(\bar{p} - \underline{q}\right) \leq \underline{p} - \underline{q} \tag{25}$$

$$(1-\theta)(1-\gamma)(\bar{p} - \bar{q}) + \gamma(\bar{p} - q_U) \leq \underline{p} - q_U \tag{26}$$

then there is a perfect Bayesian equilibrium with a type 1 strategy profile.

Proof. See appendix on page 29. □

Low price equilibria exhibit full trade as a transaction always takes place since $\beta(t_b, \rho(t_s)) =$

1 for all $(t_b, t_s) \in T$, so the buyer is always willing to buy at the price the seller offers under any possible type profile. These equilibria are said to exhibit price signaling as prices carry informative signals about quality to uninformed buyers since $E\{q|\bar{p}\} > E\{q|\underline{p}\}$, so the expected quality given a high price is higher than the expected quality given a low price. Such equilibria support bargains since an uninformed seller sets a posted price below her own value for a high quality item. In this case, an informed buyer may be able to obtain a high quality item at a bargain price.

2.3.5 High Price Equilibria

Proposition 3 states that a broad class of environments support high price equilibria with type 2 strategy profiles. Equation (27) implies the individual rationality constraints that an informed buyer and an informed seller each guarantee themselves a non-negative payoff when the item quality is high. Equation (28) implies the individual rationality constraints that an informed buyer and an informed seller each guarantee themselves a non-negative payoff when the item quality is low. It also implies that an informed high quality seller can never make a profit by selling at the low price. Equation (29) implies that, in equilibrium, an uninformed buyer is unwilling to pay the high price. Equation (30) implies that an uninformed seller is better off selling at the high price only when the item quality is high and the buyer is informed than selling at the low price with certainty.

Proposition 3. *If (ρ, β) is a type 2 strategy profile such that*

$$\bar{q} \leq \bar{p} \leq k\bar{q} \tag{27}$$

$$\underline{q} \leq \underline{p} \leq k\underline{q} \leq \bar{q} \tag{28}$$

$$k\left(\theta\lambda\underline{q} + (1-\theta)\,\bar{q}\right) \leq \bar{p}\,(\theta\lambda + 1 - \theta) \tag{29}$$

$$\theta\,(1-\gamma)\left(\underline{p} - \underline{q}\right) + \gamma\left(\underline{p} - q_U\right) \leq (1-\theta)\,(1-\gamma)\,(\bar{p} - \bar{q}) \tag{30}$$

then there is a perfect Bayesian equilibrium with a type 2 strategy profile.

Proof. See appendix on page 30. □

High price equilibria do not exhibit full trade as a transaction does not take place if an uninformed buyer meets an uninformed seller since $\beta\,(U_b, \rho\,(U_s)) = 0$. High price equilibria exhibit price signaling as prices carry informative signals about quality to uninformed buyers since $E\{q|\bar{p}\} > E\{q|\underline{p}\}$. High price equilibria do not support bargains since the buyer can never purchase an item at a price below the seller's own valuation for the item. Such equilibria are said to exhibit price signaling as informed sellers with low quality items set a lower posted price than other sellers, so uninformed buyers can infer some information about quality from observing posted prices.

2.4 Many Buyers and Sellers

Consider a market populated by n sellers of each type (high, low, and uninformed) and n buyers of each type (high, low, and uninformed). Let $t_i \in T_s$ denote the type of seller i and let $t_j \in T_b$ denote the type of buyer j. Let $\phi(q, t_i, t_j)$ denote the number of items with quality q offered by a seller of type t_i to a buyer of type t_j. Buyers value items with quality q at kq where $k > 1$. The number of low quality items offered by seller i to buyer j is given by

			t_j	
$\phi(\underline{q}, t_i, t_j)$		H_b	U_b	L_b
	H_s	0	0	0
t_i	U_s	0	$m\theta\lambda\gamma$	$m\theta\lambda(1-\gamma)$
	L_s	0	$m\theta(1-\lambda)\gamma$	$m\theta(1-\lambda)(1-\gamma)$

where $m \in \mathbb{N}$. The number of high quality items offered by seller i to buyer j is given by

			t_i	
$\phi(\bar{q}, t_i, t_j)$		H_b	U_b	L_b
	H_s	$m(1-\theta)(1-\lambda)(1-\gamma)$	$m(1-\theta)\lambda(1-\gamma)$	0
t_j	U_s	$m(1-\theta)\lambda(1-\gamma)$	$m(1-\theta)\lambda\gamma$	0
	L_s	0	0	0

Each seller i chooses a single posted price $p_i \in \mathbb{R}_+$. After observing prices, buyer j decides which items to purchase. Let $B_{ij} = 1$ if buyer j purchases the items offered by seller i. Otherwise, let $B_{ij} = 0$. Buyer j's strategy is given by $\beta_j : \mathbb{R}_+ \to \{0, 1\}$ such that $B_{ij} = \beta_j(p_i)$. The payoff to seller i is given by

$$u_i(p, \beta) = \sum_{j=1}^{3n} \sum_{q \in Q} \phi(q, t_i, t_j)\, \beta_j(p_i)(p_i - q) \tag{31}$$

Here, the payoff of a seller who faces a population of buyers employing distinct pure strategies is equivalent to the expected payoff of a seller who faces a single buyer employing a mixed strategy. The payoff of buyer j is given by

$$u_j(p, \beta) = \sum_{i=1}^{3n} \sum_{q \in Q} \phi(q, t_i, t_j) \beta_j(p_i)(kq - p_i) \tag{32}$$

Here, the payoff of a buyer who faces a population of sellers employing distinct pure strategies is equivalent to the expected payoff of a buyer who faces a single seller employing a mixed strategy. Proposition 4 states that Perfect Bayesian equilibria of the two agent interaction correspond with Nash equilibria of markets with many buyers and sellers.

Proposition 4. *If (ρ_0, β_0, μ_0) is a perfect Bayesian equilibrium of the interaction between a single buyer and a single seller then (p, β) such that $p_i = \rho_0(t_i)$ and $\beta_j(x) = \beta_0(t_j, x)$ for $i, j \in \{1, \ldots, 3n\}$ is a Nash equilibrium of the market with multiple buyers and sellers.*

Proof. See appendix on page 31. □

2.5 The Best Response Dynamic

Consider an environment where a population of buyers and a population of sellers repeatedly interact as described in section 2.4. In every period, each seller i selects a posted price p_i and each buyer selects a reservation price R_j. Buyer j accepts offers with posted prices less than or equal to her reservation price such that

$$\beta_j(p) = \begin{cases} 1 & \text{if } p \leq R_j \\ 0 & \text{if } p > R_j \end{cases}$$

The best response dynamic is an adaptive model under which agents asynchronously switch to myopic best responses. Let $\eta_i \in [0, 1]$ denote the rate at which agent i adjusts her strategy. If the current strategy profile is given by $s = (s_1, \ldots, s_n)$ then the probability that an agent

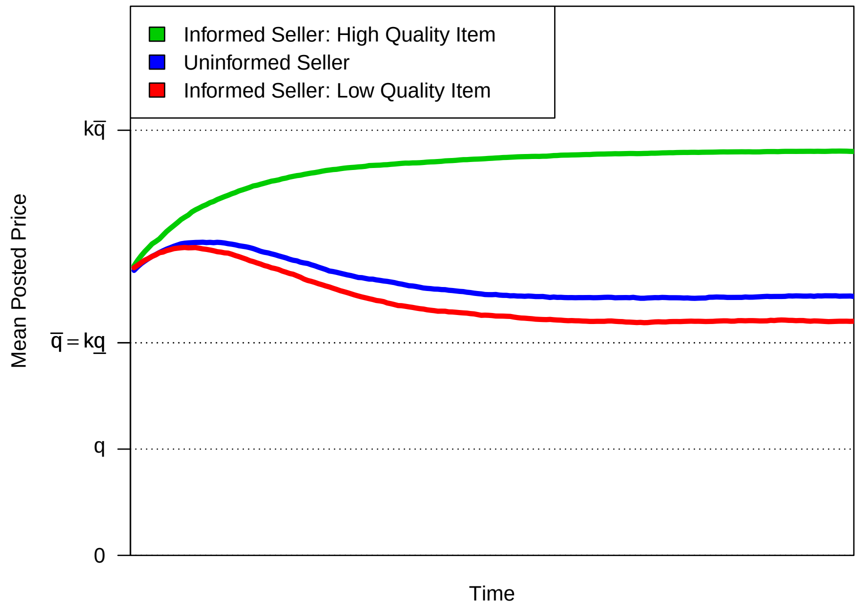

Figure 1: Mean price paths under the best response dynamic for low price equilibrium parameters $\underline{q} = 3$, $\bar{q} = 6$, $k = 2$, $\theta = \frac{7}{8}$, $\gamma = \frac{1}{8}$, $\lambda = \frac{1}{2}$.

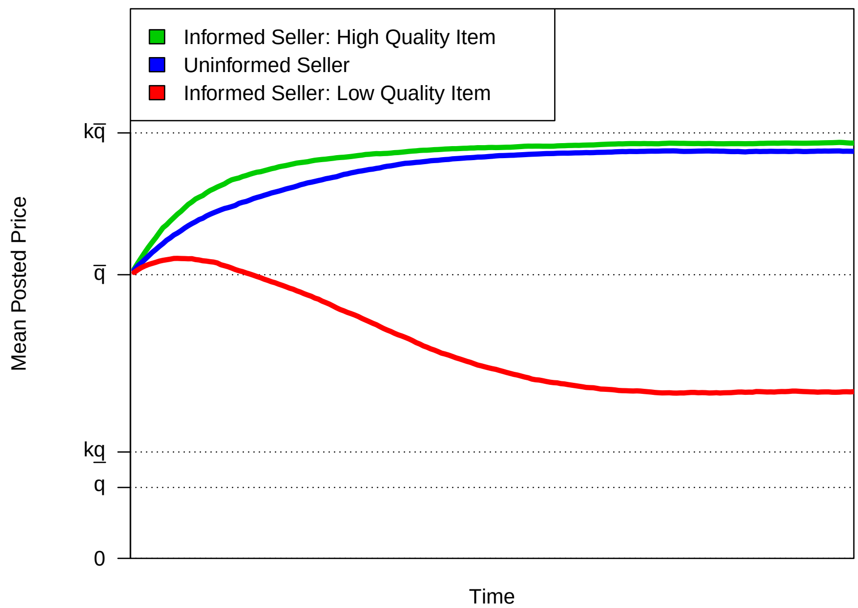

Figure 2: Mean price paths under the best response dynamic for high price equilibrium parameters $\underline{q} = 2$, $\bar{q} = 8$, $k = 1.5$, $\theta = \frac{1}{2}$, $\gamma = \frac{1}{2}$, $\lambda = \frac{1}{2}$.

i switches from her current strategy s_i to the alternate strategy s_i' is given by

$$\mathbb{P}_i\left(s_i'|s\right) = \frac{\eta_i a_i\left(s_i'|s\right)}{\sum_{x_i \in S_i} a_i\left(x_i|s\right)}$$

$$a_i\left(x_i|s\right) = \begin{cases} 1 & \text{if } x_i \in \operatorname*{argmax}_{y_i \in S_i} \pi_i\left(y_i, s_{-i}\right) \\ 0 & \text{otherwise} \end{cases}$$

Figure 1 illustrates the mean price paths predicted by the best response dynamic for $\underline{q} = 3$, $\bar{q} = 6$, $k = 2$, $\theta = \frac{7}{8}$, $\gamma = \frac{1}{8}$, $\lambda = \frac{1}{2}$. By proposition 2 these parameters are sufficient for a low price equilibrium under which informed sellers with high quality items post a high price $\bar{p} \in [\bar{q}, k\bar{q}]$ while uninformed sellers and informed sellers with low quality items post a low price $\underline{p} \in \left[\underline{q}, k\underline{q}\right]$. Consistent with these equilibrium predictions, the best response dynamic predicts that uninformed sellers and informed sellers with low type items will post similar prices.

Figure 2 illustrates the mean price paths predicted by the best response dynamic for $\underline{q} = 2$, $\bar{q} = 8$, $k = 1.5$, $\theta = \frac{1}{2}$, $\gamma = \frac{1}{2}$, $\lambda = \frac{1}{2}$. By proposition 3 these parameters are sufficient for the existence a high price equilibrium under which uninformed sellers and informed sellers with high quality items post a high price $\bar{p} \in [\bar{q}, k\bar{q}]$ while informed sellers with low quality items post a low price $\underline{p} \in \left[\underline{q}, k\underline{q}\right]$. Consistent with these equilibrium predictions, the best response dynamic predicts that uninformed sellers and informed sellers with high type items will post similar prices.

2.6 Noisy Best Response Dynamics

The noisy best response dynamic is an adaptive model under which agents asynchronously switch to noisy approximations of their myopic best responses. Let $\eta_i \in [0, 1]$ denote the rate at which agent i adjusts her strategy. Let α_i denote agent i's precision in selecting a best response. Let β_i denote agent i's sensitivity to differences in the payoffs yielded by distinct strategies. If the current strategy profile is given by $s = (s_1, \ldots, s_n)$ then the probability that an agent i switches from her current strategy s_i to an alternate strategy s_i' is given by

$$
\begin{aligned}
\mathbb{P}_i\left(s_i'|s\right) &= \frac{\eta_i \exp u_i\left(s_i'|s\right)}{\sum\limits_{x_i \in S_i} \exp u_i\left(x_i|s\right)} && (33)\\
u_i\left(x_i|s\right) &= \alpha_i \min\left|x_i - BR_i\left(s\right)\right| + \beta_i \pi_i\left(x_i, s_{-i}\right) && (34)\\
BR_i\left(s\right) &= \underset{y_i \in S_i}{\operatorname{argmax}}\ \pi_i\left(y_i, s_{-i}\right) && (35)
\end{aligned}
$$

Such agents may exhibit two distinct types of behavioral noise. The parameter α_i indexes the extent to which agent i exhibits an imprecise "trembling hand" such that she may not always select a best response but she is more likely to select strategies that are near a best response. The parameter β_i indexes the extent to which agent i exhibits a logit quantal response such that she may not always select a best response but she is more likely to select strategies that yield higher payoffs. In the limit as $\alpha_i \to 0$ and $\beta_i \to 0$ the noisy best response converges to a pure noise uniform distribution. Conversely, in the limit as $\alpha_i \to \infty$ or $\beta_i \to \infty$ the noisy best response converges to the exact best response considered in 2.5.

In the absence of behavioral noise, exact selection of a best response and exact payoff maximization are equivalent. However, in the presence of behavioral noise these two models can make different predictions. For example, relatively flat payoff functions where the best response is only slightly more profitable than other strategies will lead to greater variation in behavior when a subjects is more likely to select strategies with higher payoffs and she exhibits imperfect sensitivity to payoff differences. In contrast, if a subject is simply more likely to select strategies near her best response, then even if she exhibits imperfect precision in her strategy selection, the shape of the payoff function away from the best response does not affect her behavior.

3 Experimental Design and Procedures

This study investigates markets with many buyers and sellers under two experimental treatment conditions. Treatment 1 implements parameter values under which uninformed sellers emulate informed sellers with low quality items in equilibrium as shown by proposition 2. Treatment 2 implements parameter values under which uninformed sellers emulate informed sellers with high quality items in equilibrium as shown by proposition 3. Table 1 presents the parameter values implemented by each experimental treatment.

A total of ten experimental sessions were conducted, five for each of the two experimental

	Treatment 1	Treatment 2
High Quality ($\bar{q}$)	6	8
Low Quality ($\underline{q}$)	3	2
Gains From Trade (k)	2	3/2
Proportion Low Quality (θ)	7/8	1/2
Proportion Informed Buyers (γ)	1/8	1/2
Proportion Informed Sellers (λ)	1/2	1/2

Table 1: Experimental Treatments

treatment conditions. Each experimental session was conducted with 24 subjects, for a total of 240 experimental subjects. The experimental design was between subjects such that each session implemented only one treatment and each subject participated in only one session. All ten sessions were conducted at the Chapman University Economic Science Institute Laboratory.

During each session, 12 subjects took the role of buyers and 12 subjects took the role of sellers. Each experimental session consisted of 100 periods, each of which implemented a market with many buyers and many sellers as detailed in Section 2.4. During each period, some sellers were informed about the quality of the items while others remained uninformed. Similarly, some buyers were informed about the quality of the items they were offered while others remained uninformed. During each period each seller selected a posted price for the items they offered and each buyer selected a reservation price for the items they were offered. At the end of each period transactions took place at the posted price whenever it was below the corresponding reservation price.

Figure 3 depicts the experimental interface. Throughout each session, subjects could observe the current period, their payoff from the previous period, their endowment, their valuations for each type of item, the valuations of others, and their counterfactual payoffs from the previous period. Sellers could observe the number of offers at each quality level they made to each buyer. Buyers could observe the number of offers at each quality level they received from each seller. Similar interfaces providing counterfactual payoffs have been employed in the previous literature including Cason et al. (2013), Oprea et al. (2011), and Stephenson (2019). At the end of each session, subjects received their average payoff over all periods plus a seven dollar show up bonus with an average final payment of $18.59.

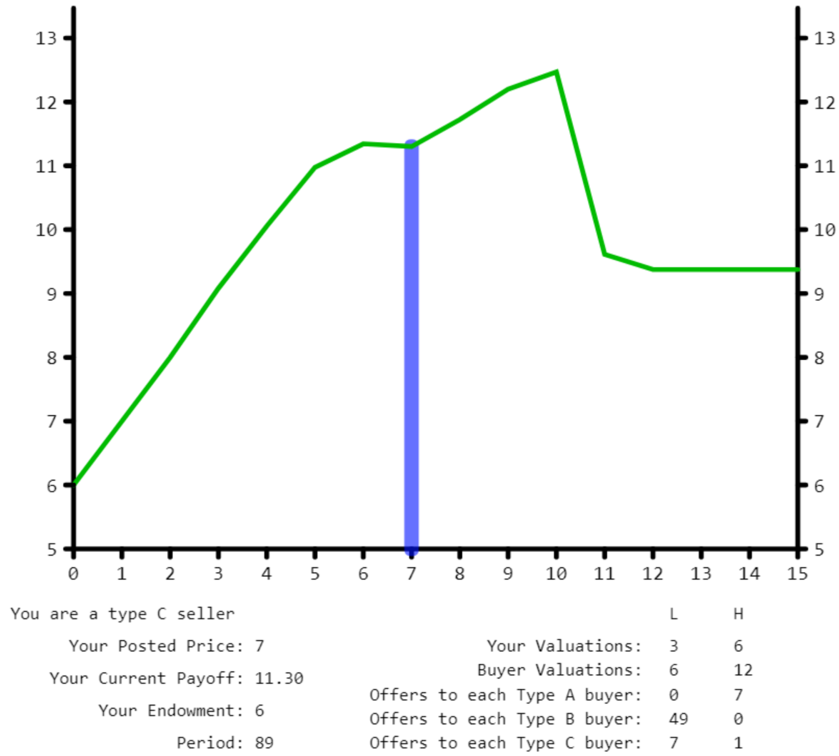

Figure 3: Screenshot of the Experimental Interface.

4 Hypotheses

Proposition 2 provides sufficient conditions for the existence of a low price equilibrium where uninformed sellers post the same low price as informed low quality sellers. These conditions are satisfied by treatment 1 of the experimental design.

Hypothesis 1. *Uninformed sellers will pool with low quality informed sellers under Treatment 1.*

Proposition 3 provides sufficient conditions for the existence of a high price equilibrium where uninformed sellers post the same high price as informed high quality sellers. These conditions are satisfied by treatment 2 of the experimental design.

Hypothesis 2. *Uninformed sellers will pool with high quality informed sellers under treatment 2.*

Under both types of equilibria, uninformed buyers can make valid inferences from prices about the quality of items they are offered. Under low price equilibria, uninformed buyers can reliably infer high quality from high prices. Under high price equilibria, uninformed buyers can reliably infer low quality from low prices.

Hypothesis 3. *Both treatments will exhibit significant price signaling.*

Low price equilibria exhibit full trade as every possible interaction between a buyer and a seller results in a transaction. In constrast, high price equilibria do not exhibit full trade since neither uninformed buyers nor informed low quality buyers are willing to pay the prices posted by uninformed sellers. Hence such interactions do not result in transactions under high price equilibria.

Hypothesis 4. *Treatment 1 will exhibit a significantly higher transaction rate than treatment 2.*

In low price equilibria, informed buyers have an opportunity to find bargains where they can purchase high quality items from uninformed sellers at prices that are lower than the seller's value for high quality items. Conversely, no such opportunities exist under high price equilibria since uninformed sellers post prices that are greater than their value for high quality items.

Hypothesis 5. *Buyers will encounter significantly more bargains in treatment 1 than treatment 2.*

5 Results

Figure 4 illustrates the mean price paths observed under treatment 1 and Figure 5 illustrates the mean price paths observed under treatment 2. In both figures, the horizontal axis indicates periods ranging from 1 to 100 and the vertical axis indicates the mean posted price over all of the sessions that implemented the respective experimental treatment. The green line illustrates the prices posted by informed sellers with high quality items. The blue line illustrates the prices posted by uninformed sellers. The red line illustrates the prices posted by informed sellers with low quality items.

Result 1. *The prices posted by uninformed sellers were similar to prices posted by low quality informed sellers under treatment 1 and were similar to prices posted by high quality informed sellers under treatment 2.*

Consistent with equilibrium predictions, the prices posted by uninformed sellers in treatment 1 are closely aligned with the prices posted by low quality informed sellers while the prices

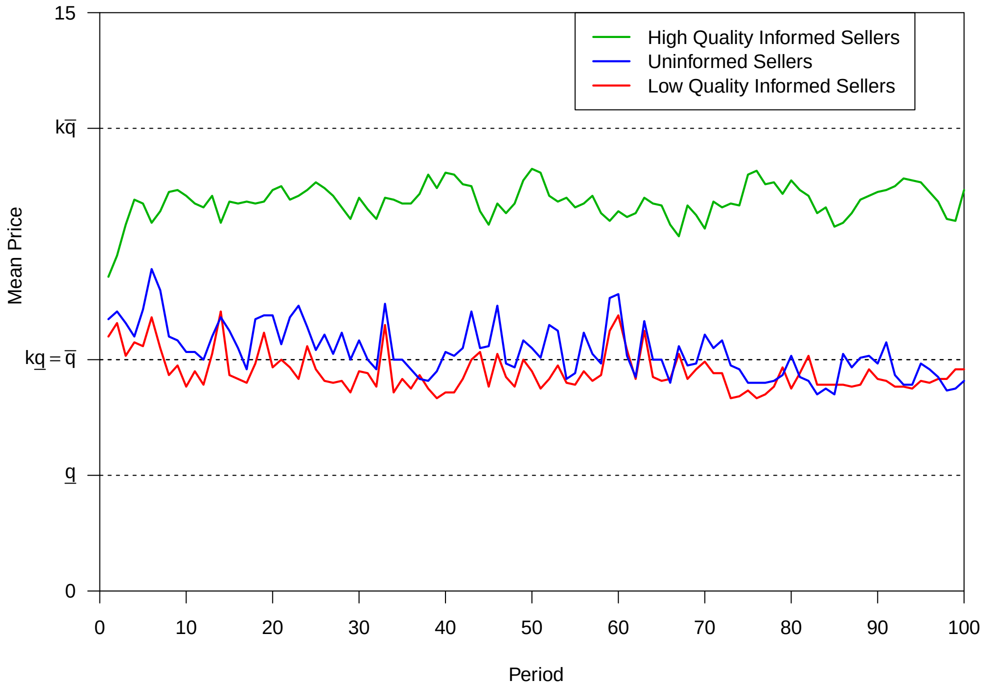

Figure 4: Mean price paths for treatment 1 (low price equilibrium)

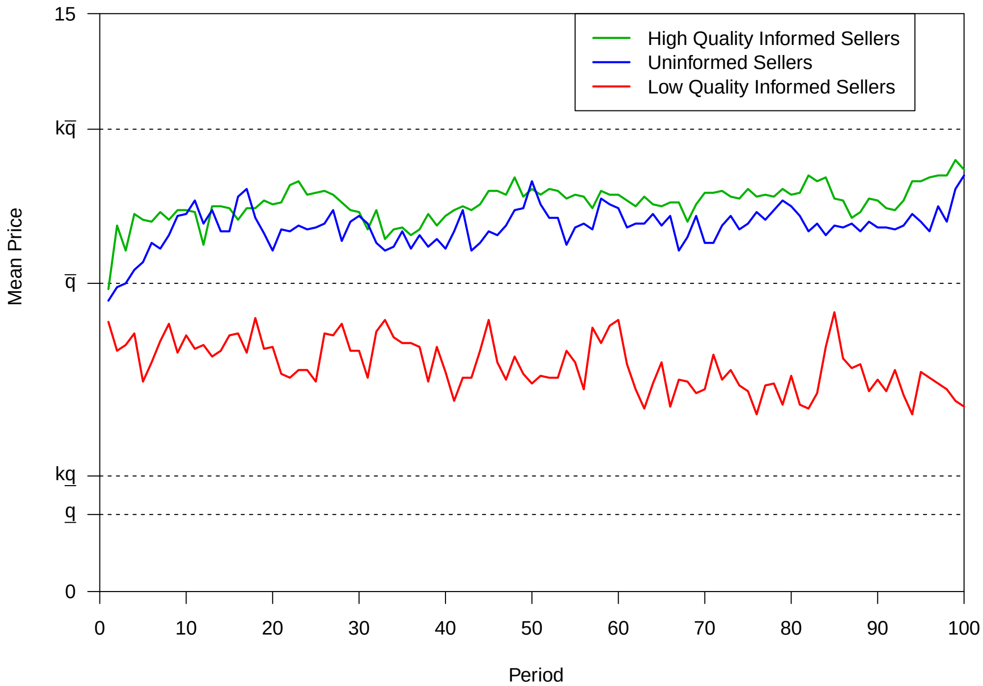

Figure 5: Mean price paths for treatment 2 (high price equilibrium)

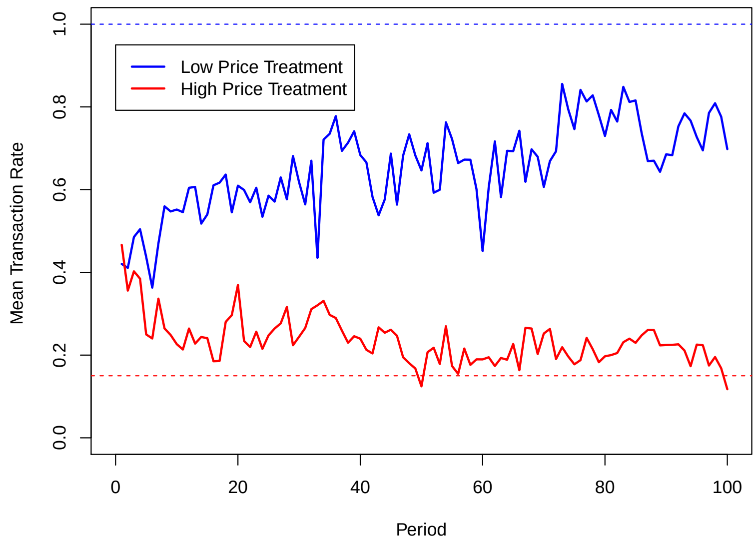

Figure 6: Mean transaction rates by period. Solid lines indicate observed fraction of items sold. Dashed lines indicate equilibrium predictions.

	Treatment		Rank-Sum Test	t-test
	1	2	p-value	p-value
Transaction Rate	0.748	0.262	0.007937	0.000003

Table 2: Hypothesis tests for differences in transaction rates across treatments. The unit of observation is one session for a total of 10 observations.

posted by uninformed sellers in treatment 2 are closely aligned with the prices posted by high quality informed sellers. Table 3 presents hypothesis tests for pooling behavior under each treatment condition. Both a non-parametric Mann-Whitney-Wilcoxon rank-sum test and a parametric t-test find that the prices posted by uninformed sellers were significantly closer to prices posted by low quality informed sellers in treatment 1 than in treatment 2. Both tests also find that the prices posted by uninformed sellers were significantly closer to prices posted by high quality informed sellers in treatment 2 than in treatment 1. These results support the pooling behavior predicted by proposition 2 and proposition 3 respectively.

The empirical price paths shown in figure 4 and figure 5 exhibit remarkable similarity with the theoretical predictions of the best response dynamics as illustrated by figure 1 and figure 2 respectively. In treatment 2, the mean posted price path for uninformed sellers and informed high quality sellers lies in the equilibrium range $[\bar{q}, k\bar{q}]$. However, as predicted by the best response dynamic, the mean posted price path for informed sellers and low quality sellers in treatment 2 lies above the equilibrium range $\left[\underline{q}, k\underline{q}\right]$.

Result 2. *Prices carried significant information about quality under both treatments.*

	Treatment 1	Treatment 2	Rank-Sum Test p-value	t-test p-value
$\lvert P_{\text{Uninformed}} - P_{\text{LowQuality}}\rvert$	\$0.74	\$3.59	0.007937	0.00239
$\lvert P_{\text{Uninformed}} - P_{\text{HighQuality}}\rvert$	\$3.88	\$0.62	0.007937	0.00002

Table 3: Hypothesis tests regarding pooling behavior. The unit of observation is one session for a total of 10 observations.

	Informed Seller Type		Rank-Sum Test	t-test
Price	High Quality	Low Quality	p-value	p-value
Treatment 1	\$10.10	\$5.69	0.007937	0.000001
Treatment 2	\$10.08	\$5.88	0.007937	0.000065

Table 4: Hypothesis tests for differences in posted prices across seller types. The unit of observation is the mean posted price in one session of one type of seller for a total of 10 observations.

	Treatment 1	Treatment 2	Rank-Sum Test p-value	t-test p-value
Bargain Rate	0.4815	0.1715	0.01587	0.000997

Table 5: Hypothesis tests for differences in the proportion of bargains offered by uninformed sellers across treatments. The unit of observation is one session for a total of 10 observations.

Table 4 presents hypothesis tests for differences in posted prices across item qualities. Nonparametric Mann-Whitney-Wilcoxon rank-sum tests and parametric t-tests find that informed high quality sellers posted significantly higher prices than informed low quality sellers in both treatments. Consequently, uninformed buyers could make valid inferences about item qualities based on the posted prices. Consistent with equilibrium predictions, both treatments exhibited significant price signaling.

Figure 6 illustrates the mean transaction rates observed under each treatment. The vertical axis indicates the mean fraction of items sold. The horizontal axis indicates periods ranging from 1 to 100. The solid blue line indicates the mean fraction of items sold over all of the sessions that implemented treatment 1. The solid red line indicates the mean fraction of items sold under all sessions that implemented treatment 2. The dotted lines indicate equilibrium predictions. In the first period, both treatments exhibited similar transaction rates. In later periods, treatment 1 consistently exhibited higher transaction rates than treatment 2.

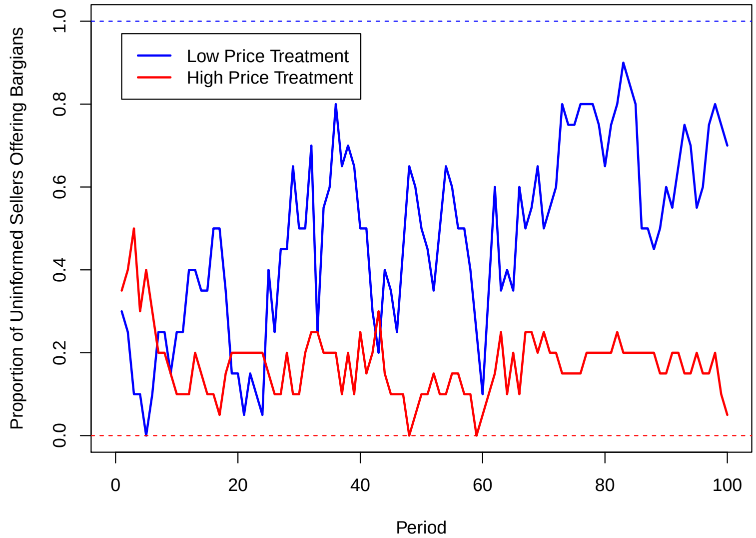

Figure 7: Proportion of uninformed sellers offering high quality items at prices less than the seller's own valuation for high quality items. Solid lines indicate observed fraction of items sold. Dashed lines indicate equilibrium predictions.

Result 3. *Observed transaction rates were significantly higher in treatment 1 than treatment 2.*

Table 2 presents hypothesis tests for differences in the observed transaction rate across experimental treatment conditions. Both a non-parametric Mann-Whitney-Wilcoxon rank-sum test and a parametric t-test find that the observed transaction rate in treatment 1 was significantly higher than the observed transaction rate in treatment 2.

Figure 7 illustrates the proportion of uninformed sellers offering bargains in each treatment. Uninformed sellers are said to offer bargains if their posted prices for a high quality item are less than their own valuation for high quality items. The horizontal axis indicates periods ranging from 1 to 100. Solid lines indicate proportion of uninformed sellers offering bargains in a given period. The dotted lines indicate equilibrium predictions. In the first fifty periods, this bargain rate was highly volatile. Over the last fifty periods, this bargain rate was consistently higher in treatment 1 than in treatment 2.

Result 4. *Significantly more bargains were available under treatment 1 than under treatment 2.*

Table 5 presents hypothesis tests for differences in the proportion of uninformed sellers offering bargains between experimental treatment conditions. Both a non-parametric Mann-Whitney-Wilcoxon rank-sum test and a parametric t-test find that the fraction of uninformed sellers offering bargains in treatment 1 was significantly higher than in treatment 2.

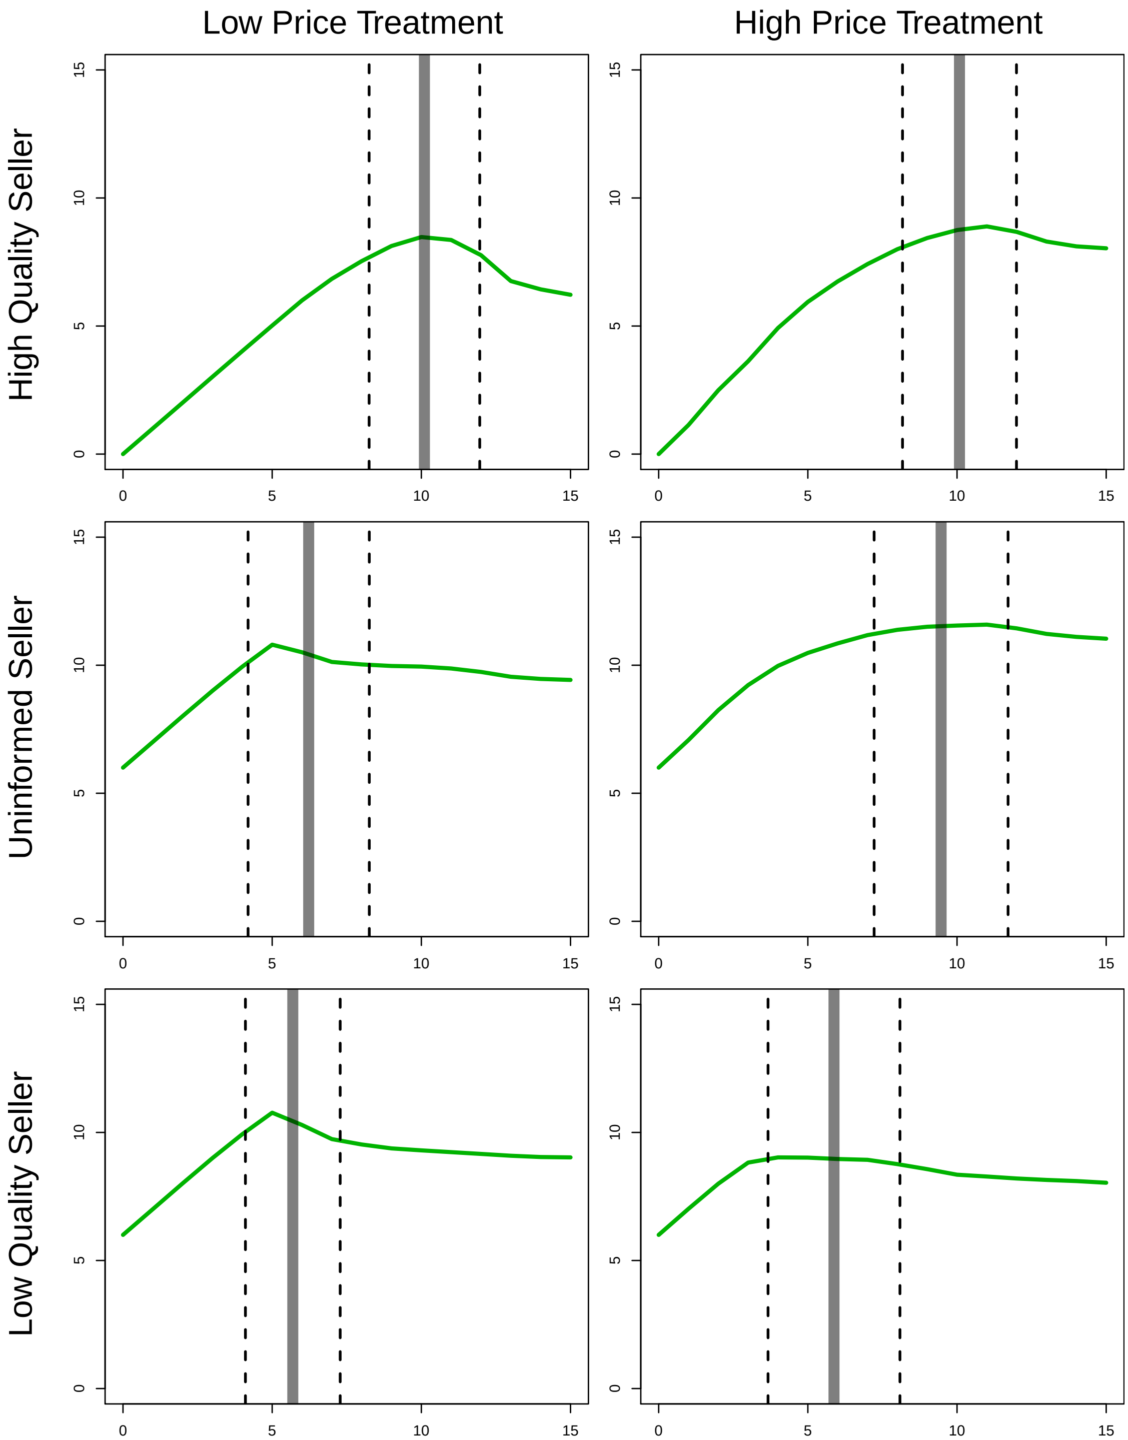

Figure 8: Average posted prices and payoffs by seller type and treatment over all 100 periods. The horizontal axis indicates posted prices and the vertical axis indicates payoffs. The green line indicates the average payoff to each possible posted price. The solid gray line indicates the mean posted price. The dotted gray lines indicate one standard deviation from the mean.

Figure 8 illustrates the average posted prices and payoffs by seller type over all 100 periods of each treatment. The left column illustrates the observed posted prices for each type of seller under treatment 1. The right column illustrates the observed posted prices for each type of seller under treatment 2. The horizontal axis indicates posted prices and the vertical axis indicates payoffs. The green line indicates the average payoff to each possible posted price. The solid gray line indicates the mean posted price. The dotted gray lines indicate one standard deviation from the mean.

Under both treatments, the mean posted price for each type of seller is nearly optimal, indicating that subjects responded to incentives. Yet posted prices also exhibit considerable variance, suggesting the presence of individual heterogeneity or behavioral noise. To investigate individual level behavior we calculate the maximum likelihood estimates for the the noisy best response model described in Section 2.6 where each subject exhibits a distinct level of precision in her selection of optimal strategies and a distinct level of sensitivity to payoff differences.

Figure 9 illustrates each subject's maximum likelihood parameter estimates. Each point indicates the parameters estimated for a single subject. The vertical axis illustrates a given subject's estimated level of sensitivity to payoff differences. The horizontal axis illustrates a given subject's tendency to select strategies near her best response. Some subjects exhibited strong sensitivity to payoff differences but relatively little tendency to select strategies near their best response. Other subjects exhibited a strong tendency to select strategies near their best response but relatively weak sensitivity to payoff differences.

In the absence of behavioral noise, exact selection of a best response and exact payoff maximization are equivalent. However, in the presence of behavioral noise these two models can make different predictions. For example, relatively flat payoff functions where the best response is only slightly more profitable than other strategies will lead to greater variation in behavior when a subjects is more likely to select strategies with higher payoffs and she exhibits imperfect sensitivity to payoff differences. In contrast, if a subject is simply more likely to select strategies near her best response, then even if she exhibits imperfect precision in her strategy selection, the shape of the payoff function away from the best response does not affect her behavior.

Result 5. *Subjects exhibited heterogeneous payoff sensitivity and strategy precision.*

Figure 10 illustrates the p-values for likelihood ratio tests for restrictions to a single source of behavioral noise. Each point indicates the p-values obtained for a single subject. The vertical axis displays the p-value for a test of the null hypothesis that imprecision in strategy

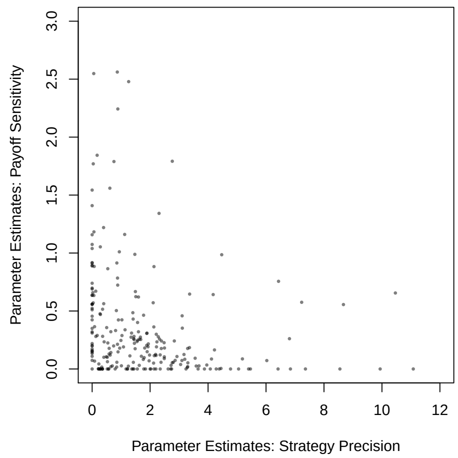

Figure 9: Maximum likelihood parameter estimates for payoff sensitivity and price precision. Each point indicates the parameters estimated for a single subject.

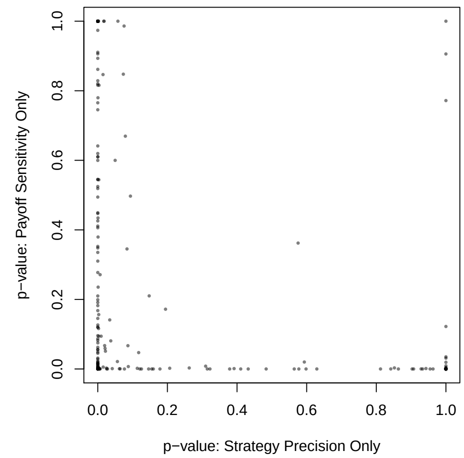

Figure 10: Likelihood ratio test p-values for restrictions to a single source of behavioral noise. Each point indicates the p-values obtained for a single subject.

selection is the only source of noise in a given subject's behavior. The horizontal axis displays the p-value for a test of the null hypothesis that imperfect sensitivity to payoff differences is the only source of noise in a given subject's behavior. For many subjects we can reject exactly one of these hypotheses at the 1% level, indicating that some subjects focused on selecting strategies near their best response while others focused on selecting strategies with relatively large payoffs.

6 Conclusion

Akerlof's (1970) classic analysis of lemons markets assumes all sellers are informed and all buyers are uninformed about the quality of goods for sale. In this paper, we consider a more general environment in which some buyers are informed and some sellers are informed. In this environment, we characterize two types of perfect Bayesian Nash equilibria in which sellers endogenously set prices and buyers endogenously form beliefs.

Low price equilibria exhibit full trade while high price equilibria exhibit only partial trade. Bargains are offered by uninformed sellers under low price equilibria but not under high price equilibria. Prices serve as reliable signals of quality to uninformed buyers under both types of equilibria. These theoretical predictions suggest Pareto efficient full trade, price signaling, and bargains can all coexist in markets with asymmetric information.

While subject behavior was largely consistent with the equilibrium pooling predictions, it systematically deviated from equilibrium predictions about the transaction rate and the bargain rate. Equilibrium predicts full trade in the first treatment but low levels of trade in the second treatment. While the first treatment exhibited significantly higher transaction rates, it remained significantly below the equilibrium predictions. Similarly, equilibrium predicts that uninformed sellers will always offer bargains in our first treatment but will never offer bargains in our second treatments. While the first treatment exhibited significantly higher bargain rates, it remained significantly below equilibrium predictions.

We found that subjects exhibited two distinct types of behavioral noise in their selection of strategies. Some subjects focused on selecting strategies near a best response, while other subjects focused on selecting strategies with higher payoffs. These two methods of optimization are indistinguishable in the absence of behavioral noise, since exact selection of a best response is equivalent to exact payoff maximization. Only in the presence of behavioral noise do these two methods yield distinct behavior.

In contrast with conventional adverse selection models, our results indicate that price can serve as a reliable signal of quality in markets with asymmetric information. Casual observation suggests that consumers often try to infer product quality from prices. For instance, when deciding on which wine to purchase, an uninformed consumer might use the price as a proxy for quality. In our model, this arises because the presence of some informed buyers can incentivize informed sellers to charge prices that reflect item quality, allowing the uninformed buyers to free ride on this price signal. Future research could generalize our model to the case of continuous quality distributions and identify the degree to which our results extend to naturally occurring markets with asymmetric information.

References

George A Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

Kyle Bagwell and Michael H Riordan. High and declining prices signal product quality. *The American Economic Review*, pages 224–239, 1991.

Timothy N Cason, Daniel Friedman, and Ed Hopkins. Cycles and instability in a rock–paper–scissors population game: A continuous time experiment. *Review of Economic Studies*, 81(1):112–136, 2013.

Alecia Cassidy. How does mandatory energy efficiency disclosure affect housing prices? *Available at SSRN 3047417*, 2019.

Yuk-Shee Chan and Hayne Leland. Prices and qualities in markets with costly information. *The Review of Economic Studies*, 49(4):499–516, 1982.

Russell Cooper and Thomas W Ross. Prices, product qualities and asymmetric information: The competitive case. *The Review of Economic Studies*, 51(2):197–207, 1984.

Giuseppe Dari-Mattiacci, Sander Onderstal, and Francesco Parisi. Inverse adverse selection: The market for gems. *Minnesota Legal Studies Research Paper*, (10-47):2010–04, 2011.

Wenyan Hu and Alvaro Bolivar. Online auctions efficiency: a survey of ebay auctions. In *Proceedings of the 17th international conference on World Wide Web*, pages 925–934. ACM, 2008.

Maarten CW Janssen and Santanu Roy. Signaling quality through prices in an oligopoly. *Games and Economic Behavior*, 68(1):192–207, 2010.

Anke S Kessler. Revisiting the lemons market. *International Economic Review*, 42(1):25–41, 2001.

Erica Myers, Steven Puller, and Jeremy West. Effects of mandatory energy efficiency disclosure in housing markets. 2019.

Ryan Oprea, Keith Henwood, and Daniel Friedman. Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6): 2206–2225, 2011.

Steven Salop and Joseph Stiglitz. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, 44(3):493–510, 1977.

Daniel Stephenson. Coordination and evolutionary dynamics: When are evolutionary models reliable? *Games and Economic Behavior*, 113:381–395, 2019.

Asher Wolinsky. Prices as signals of product quality. *The Review of Economic Studies*, 50 (4):647–658, 1983.

A Proofs

Proof of proposition 1. If $\rho(L_s) > \rho(H_s)$ then $\beta(H_b, \rho(L_s)) = \beta(L_b, \rho(L_s)) = 1$ since $E\{q|H_b\} = \bar{q} \geq \underline{q} = E\{q|L_b\}$. But then the informed high quality seller could earn a higher expected profit by mimicking the informed low quality seller. So we must have $\rho(L_s) \leq \rho(H_s)$.

If $\beta(L_b, \rho(U_s)) = 0$ and $\rho(L_s) < \rho(U_s) < \rho(H_s)$ then $\beta(U_b, \rho(H_s)) = \beta(H_b, \rho(H_s)) = 1$ since $E\{q|p = \rho(H_s)\} = \bar{q}$. But then the uninformed seller could earn a higher expected profit by mimicking the informed high quality seller.

If $\beta(L_b, \rho(U_s)) = 1$ and $\rho(L_s) < \rho(U_s) < \rho(H_s)$ then $\beta(U_b, \rho(U_s)) = 1$ since $E\{q|U_b, p = \rho(U_s)\} \geq E\{q|L_b, p = \rho(U_s)\}$. But then the low quality informed seller could earn a higher expected profit by mimicking the uninformed seller.

If $\rho(U_s) < \rho(L_s)$ then $\beta(U_b, \rho(L_s)) = \beta(L_b, \rho(L_s)) = 1$ since $E\{q|p = \rho(L_s)\} \geq \underline{q}$. But then the uninformed seller could earn a higher expected profit by mimicking the informed low quality seller.

If $\rho(U_s) > \rho(H_s)$ and $\beta(U_b, \rho(U_s)) = 1$ then $\beta(H_b, \rho(U_s)) = 1$ since $E\{q|H_b\} \geq E\{q|U_b\}$. But then the informed high quality seller could earn a higher expected profit by mimicking the uninformed seller.

If $\rho(U_s) > \rho(H_s)$ and $\beta(U_b, \rho(U_s)) = 0$ then $\beta(L_b, \rho(U_s)) = 0$ by the incentive compatibility condition of the low type seller. Hence $\beta(H_b, \rho(U_s)) = 1$ by the incentive compatibility condition of the uninformed seller. So $\beta(U_b, \rho(H_s)) = 1$ by the incentive compatibility condition of the informed high quality seller. Then by the uninformed seller's incentive compatibility condition we have

$$\begin{aligned}
E\{\pi_s|U_s, p = \rho(U_s)\} &\geq E\{\pi_s|U_s, p = \rho(H_s)\} \\
(1-\theta)(1-\gamma)(\rho(U_s) - \bar{q}) &\geq (1-\theta)(1-\gamma)(\rho(H_s) - \bar{q}) + \gamma(\rho(H_s) - q_U) \\
(1-\theta)(1-\gamma)(\rho(U_s) - \bar{q}) &> (1-\theta)(1-\gamma)(\rho(H_s) - \bar{q}) + \gamma(\rho(H_s) - \bar{q}) \\
(1-\theta)(1-\gamma)(\rho(U_s) - \bar{q}) &> (1-\theta-\gamma+\theta\gamma)(\rho(H_s) - \bar{q}) + \gamma(\rho(H_s) - \bar{q}) \\
(1-\theta)(1-\gamma)(\rho(U_s) - \bar{q}) &> (1-\theta+\theta\gamma)(\rho(H_s) - \bar{q}) \\
(1-\theta)(1-\gamma)(\rho(U_s) - \bar{q}) &> (1-\theta)(\rho(H_s) - \bar{q}) \\
(1-\gamma)(\rho(U_s) - \bar{q}) &> \rho(H_s) - \bar{q} \\
E\{\pi_s|H_s, p = \rho(U_s)\} &> E\{\pi_s|H_s, p = \rho(H_s)\}
\end{aligned}$$

But then the informed high quality seller could earn a higher expected profit by mimicking the uninformed seller. $\square$

Proof of proposition 2. The buyer's belief μ is consistent along the path of play if $\mu\left(\underline{q}|p, L_b\right) = 1$, $\mu\left(\underline{q}|p, H_b\right) = 0$, and

$$\mu\left(\underline{q}|p, U_b\right) = \begin{cases} 1 & \text{if } p > \bar{p} \\ 0 & \text{if } p \in \left(\underline{p}, \bar{p}\right] \\ \frac{\theta}{\theta+\lambda(1-\theta)} & \text{if } p \leq \underline{p} \end{cases} \tag{36}$$

Then the buyer's expected quality satisfies

$$\underline{q} < \mathbb{E}_{\mu}\left\{q|\underline{p}, U_b\right\} = \frac{\theta \underline{q} + \lambda\left(1-\theta\right)\bar{q}}{\theta + \lambda\left(1-\theta\right)} < \bar{q} \tag{37}$$

Hence $\beta\left(t_b, p\right)$ is optimal for all t_b and p by (23) and (24). Then the seller's expected payoff is

$$\mathbb{E}\left\{\pi_s|\bar{p}, L_s\right\} = \gamma\left(\bar{p} - \underline{q}\right) \tag{38}$$

$$\mathbb{E}\left\{\pi_s|\underline{p}, L_s\right\} = \underline{p} - \underline{q} \tag{39}$$

$$\mathbb{E}\left\{\pi_s|\bar{p}, H_s\right\} = \bar{p} - \bar{q} \tag{40}$$

$$\mathbb{E}\left\{\pi_s|\underline{p}, H_s\right\} = \underline{p} - \bar{q} \tag{41}$$

$$\mathbb{E}\left\{\pi_s|\bar{p}, U_s\right\} = \left(1-\theta\right)\left(1-\gamma\right)\left(\bar{p} - \bar{q}\right) + \gamma\left(\bar{p} - q_U\right) \tag{42}$$

$$\mathbb{E}\left\{\pi_s|\underline{p}, U_s\right\} = \underline{p} - q_U \tag{43}$$

Hence $\rho\left(t_s\right)$ is optimal for all t_s by (25) and (26). $\square$

Proof of proposition 3. The buyer's belief μ is consistent along the path of play if $\mu\left(\underline{q}|p, L_b\right) = 1$, $\mu\left(\underline{q}|p, H_b\right) = 0$, and

$$\mu\left(\underline{q}|p, U_b\right) = \begin{cases} \frac{\theta\lambda}{\theta\lambda+1-\theta} & \text{if } p \in \left(\underline{p}, \bar{p}\right] \\ 1 & \text{otherwise} \end{cases} \tag{44}$$

Then the buyer's expected quality satisfies

$$\underline{q} < \mathbb{E}_b\left\{q|\bar{p}, U_b\right\} = \frac{\theta\lambda\underline{q} + (1-\theta)\,\bar{q}}{\theta\lambda + 1 - \theta} < \bar{q} \tag{45}$$

Hence $\beta\left(t_b, p\right)$ is optimal for all t_b and p by (27), (28), and (29). Then the seller's expected payoff is

$$\begin{aligned}
\mathbb{E}\left\{\pi_s|\bar{p}, L_s\right\} &= 0 && (46)\\
\mathbb{E}\left\{\pi_s|\underline{p}, L_s\right\} &= \underline{p} - \underline{q} && (47)\\
\mathbb{E}\left\{\pi_s|\bar{p}, H_s\right\} &= (1-\gamma)\left(\bar{p} - \bar{q}\right) && (48)\\
\mathbb{E}\left\{\pi_s|\underline{p}, H_s\right\} &= \underline{p} - \bar{q} && (49)\\
\mathbb{E}\left\{\pi_s|\bar{p}, U_s\right\} &= (1-\theta)(1-\gamma)\left(\bar{p} - \bar{q}\right) && (50)\\
\mathbb{E}\left\{\pi_s|\underline{p}, U_s\right\} &= \theta(1-\gamma)\left(\underline{p} - \underline{q}\right) + \gamma\left(\underline{p} - q_U\right) && (51)\\
&\quad + (1-\theta)(1-\gamma)\left(\underline{p} - \bar{q}\right)
\end{aligned}$$

Hence $\rho\left(t_s\right)$ is optimal for all $t_s \in T_s$ by (27), (28), and (30). $\square$

Proof of proposition 4. Let (ρ_1, β_1, μ_1) be a perfect Bayesian Nash equilibrium for the interaction between a single buyer and a single seller. Let (p, β) such that $p_i = \rho_1(t_i)$ and $\beta_j = \beta_1(t_b, \cdot)$ for $i, j \in \{1, \ldots, 3n\}$. Then the expected payoff to the seller in the two-agent interaction is given by

$$\begin{aligned}
\mathbb{E}\left\{\pi_s | t_s, p\right\} &= \sum_{t_b \in T_b} \mathbb{P}\left(t_b | t_s\right) \beta_0\left(t_b, p\right)\left(p - \mathbb{E}\left\{q | t_b, t_s\right\}\right) && (52)\\
&= \frac{1}{\mathbb{P}\left(t_s\right)} \sum_{t_b \in T_b} \mathbb{P}\left(t_s \cap t_b\right) \beta_0\left(t_b, p\right)\left(p - \mathbb{E}\left\{q | t_b, t_s\right\}\right) && (53)\\
&= \frac{1}{\mathbb{P}\left(t_s\right)} \sum_{t_b \in T_b} \sum_{q \in Q} \mathbb{P}\left(q \cap t_s \cap t_b\right) \beta_0\left(t_b, p\right)(p - q) && (54)\\
&= \frac{1}{m\mathbb{P}\left(t_s\right)} \sum_{t_b \in T_b} \sum_{q \in Q} \phi\left(q, t_s, t_b\right) \beta_0\left(t_b, p\right)(p - q) && (55)\\
&= \frac{1}{mn\mathbb{P}\left(t_s\right)} \sum_{j=1}^{3n} \sum_{q \in Q} \phi\left(q, t_s, t_b\right) \beta_j(p)(p - q) && (56)
\end{aligned}$$

Hence p_i maximizes $u_i(p, \beta)$. Let $H_j(p)$ denote the set of posted prices for offers to buyer j under the price profile p. The expected payoff to the buyer in the market with multiple buyers and sellers is given by

$$\begin{aligned}
u_j(p, \beta) &= \sum_{i=1}^{3n} \sum_{q \in Q} \phi\left(q, t_i, t_j\right) \beta_j\left(p_i\right)\left(kq - p_i\right) && (57)\\
&= \sum_{i=1}^{3n} \sum_{q \in Q} \phi\left(q, t_i, t_j\right) \beta_0\left(t_j, p_i\right)\left(kq - p_i\right) && (58)\\
&= n \sum_{t_s \in T_s} \sum_{q \in Q} \phi\left(q, t_s, t_j\right) \beta_0\left(t_j, \rho_0\left(t_s\right)\right)\left(kq - \rho_0\left(t_s\right)\right) && (59)\\
&= mn\mathbb{P}\left(t_b = t_j\right) \sum_{p \in H_j(\rho)} \left(k\mathbb{E}_\mu\left\{q | t_b = t_j, p\right\} - p\right) \beta_0\left(t_j, p\right) && (60)
\end{aligned}$$

Hence β_j maximizes $u_j(p, \beta)$. □