

Competing with Asking Prices*

Benjamin Lester[†] Ludo Visschers[‡] Ronald Wolthoff[§]

May 18, 2016

Abstract

In many markets, sellers advertise their good with an *asking price*. This is a price at which the seller will take his good off the market and trade immediately, though it is understood that a buyer can submit an offer below the asking price and that this offer may be accepted if the seller receives no better offers. We construct an environment with a few simple, realistic ingredients and demonstrate that, by using an asking price, sellers both maximize their revenue and implement the efficient outcome in equilibrium. We provide a complete characterization of this equilibrium and use it to explore the implications of this pricing mechanism for transaction prices and allocations.

Keywords: Asking Prices, Posted Prices, Auctions, Competing Mechanisms, Competitive Search

JEL codes: C78, D44, D82, D83, R31

*We would like to thank the editor, Johannes Hörner, three referees, Jim Albrecht, Gadi Barlevy, Jan Eeckhout, Andrea Galeotti, Pieter Gautier, Lu Han, Benoit Julien, Philipp Kircher, Jonathan Levin, Guido Menzio, Peter Norman, József Sákovics, Shouyong Shi, Andy Skrzypacz, Xianwen Shi, Rob Shimer, Susan Vroman, Gábor Virág, and Randy Wright for helpful comments. Ludo Visschers gratefully acknowledges financial support from the UK Economic and Social Research Council (ESRC), Award reference ES/N00776X/1. Ronald Wolthoff gratefully acknowledges financial support from the Connaught Fund at the University of Toronto. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

[†]Federal Reserve Bank of Philadelphia. E-mail: benjamin.lester@phil.frb.org.

[‡]University of Edinburgh, Universidad Carlos III de Madrid, & CESifo. E-mail: ludo.visschers@ed.ac.uk.

[§]University of Toronto. E-mail: ronald.p.wolthoff@gmail.com.

1 Introduction

In this paper, we consider an environment in which a trading mechanism that we call an *asking price* emerges as an optimal way of coping with certain frictions. In words, an asking price is a price at which a seller commits to taking his good off the market and trading immediately. However, it is understood that a buyer can submit an offer below the asking price and this offer could potentially be accepted if the seller receives no better offer.¹ Though asking prices are prevalent in a variety of markets, they have received relatively little attention in the academic literature. We construct an environment with a few simple, realistic ingredients and demonstrate that using an asking price is both revenue-maximizing and efficient, i.e., sellers optimally choose to use the asking price mechanism and, in equilibrium, the asking prices they select implement the solution to the planner’s problem. We provide a complete characterization of this equilibrium and explore the implications of this pricing mechanism for transaction prices and allocations.

At first glance, one might think that committing to an asking price would be sub-optimal from a seller’s point of view. After all, the seller is not only placing an upper bound on the price that a buyer might propose, he is also committing not to meet with any additional prospective buyers once the asking price has been offered. Hence, when a buyer purchases the good at the asking price, the seller has forfeited any additional rents that either this buyer or other prospective buyers were willing to pay. And yet, an asking price seems to play a prominent role in the sale of many goods (and services). Hence, a natural question is: how and why can this mechanism be optimal?

Loosely speaking, our answer requires two ingredients. First, the sellers in our environment *compete to attract buyers*. In particular, in contrast to the large literature that studies the performance of various trading mechanisms (e.g., auctions) when the number of buyers is fixed, we assume that there are many sellers, each one posts (and commits to) the process by which their good will be sold, and buyers strategically choose to visit only those sellers who promise the maximum expected payoff. Thus, in our framework, the number of buyers at each seller is endogenous and responds to the mechanism that the seller posts. The second ingredient is that the goods for sale are *inspection goods* and inspection is *costly*: though all sellers’ goods appear ex ante identical, in fact each buyer has an idiosyncratic (private) valuation for each good which can only be learned through a process of costly inspection.

Intuitively, when buyers have to incur a cost to learn their valuation and trade, a natural tension arises. *Ceteris paribus*, a seller wants to maximize the number of buyers who inspect his good and

¹What we call an asking price goes by several other names as well, including an “offering price” (as in the sale of a company), a “list price” (as in the sale of houses and cars), or a “buy-it-now” or “take-it” price (as in certain online marketplaces). In many classified advertisements, what we call an asking price often comes in the form of a price followed by the comment “or best offer.” Though the terminology may differ across these various markets, along with the fine details of how trade occurs, we think our analysis identifies an important, fundamental reason that sellers might find this basic type of pricing mechanism optimal.

the offers that these buyers make. However, without any device to limit the number of buyers or the offers they make, buyers might be hesitant to approach such a seller, as the cost of inspecting the good could outweigh the potential benefits of discovering their valuation and trying to outbid the other buyers. Instead, buyers would only incur the up-front costs of inspection if they are assured a reasonable chance of actually acquiring the good at a reasonable price.

We show that, by setting an asking price, sellers can provide buyers with this assurance. For one, an asking price implements a stopping rule, allocating the good to the first buyer who has a sufficiently high valuation, and sparing the remaining potential buyers from inspecting the good “in vain.” In this sense, the asking price mechanism in our model serves as a promise by sellers not to waste the buyers’ time and energy when they have only a small chance of actually getting the good. Moreover, by committing to sell at a particular price, a seller who uses an asking price is also assuring potential buyers that they will receive some gains from trade in the event that they discover a high valuation. For these reasons, sellers have a strong incentive to use an asking price when they are selling inspection goods.

However, in the absence of competition, sellers also have a strong incentive to charge buyers a fee to inspect their good, extracting all of the additional surplus they created by eliminating inefficient inspections. This is where our second ingredient is important: we show that, in an environment in which sellers compete to attract buyers, the optimal fee is, in fact, zero. Thus, in an environment with these two ingredients, sellers maximize profits by posting a simple mechanism, composed of one asking price for all buyers, and no fees, side payments, or other rarely observed devices. Moreover, the asking price that sellers choose ultimately maximizes the expected surplus that they create, so that equilibrium asking prices implement the planner’s solution.

Having provided the rough intuition, we now discuss our environment and main results in greater detail. As we describe explicitly in Section 2, we consider a market with a measure of sellers, each endowed with one indivisible good, and a measure of buyers who each have unit demand. Though goods appear *ex ante* identical, each buyer has an idiosyncratic valuation for each good and this valuation can only be learned through a costly inspection process. We assume that sellers have the ability to communicate *ex ante* (or “post”) how their good is going to be sold, and buyers can observe what each seller posts and visit the seller that offers the highest expected payoff. The search process, however, is frictional: each buyer can only visit a single seller and he cannot coordinate this decision with other buyers. As a result, the number of buyers to arrive at each seller is a random variable with a distribution that depends on what the seller posted.

As a first step, in Section 3 we characterize the solution to the problem of a social planner who maximizes total surplus, subject to the frictions described above—in particular, the search frictions and the requirement that a buyer’s valuation is costly to learn. The solution has three properties. First, the planner instructs buyers to randomize evenly across sellers. Second, once a

random number of buyers arrive at each seller, the planner instructs buyers to undergo the costly inspection process sequentially, preserving the option to stop after each inspection and allocate the good to one of the buyers who have learned their valuation. Finally, we characterize the optimal stopping rule for this strategy and establish that it is stationary; that is, it depends on neither the number of buyers who have inspected the good nor the realization of their valuations.

Next, we move on to our main contribution: we consider the decentralized economy where sellers use an asking price mechanism, characterize the equilibrium, and study its efficiency properties. Given the nature of the planner's optimal trading protocol, the asking price mechanism is a natural candidate to implement the efficient outcome. First, since buyers' valuations are privately observed, the asking price provides sellers a channel to elicit information about these valuations. Second, since the asking price triggers immediate trade, it implements a stopping rule, thus preventing additional buyers from incurring the inspection cost when the current buyer draws a sufficiently high valuation. Finally, since the seller also allows bids below the asking price, he retains the option to recall previous offers in which there was a positive match surplus.

These features are captured by the following game, which we study in Section 4. First, sellers post an asking price, which all buyers observe. Given these asking prices, each buyer then chooses to visit the seller (or mix between sellers) offering the maximal expected payoff. Once buyers arrive at their chosen seller, they are placed in a random order. Buyers are told neither the number of other buyers who have arrived, nor their place in the queue.² The first buyer incurs the inspection cost, learns his valuation, and can either purchase the good immediately at the asking price or submit a counteroffer. If he chooses the former, trade occurs and all remaining buyers at that particular seller neither inspect the good nor consume. If he chooses the latter, the seller moves on to the second buyer (if there is one) and the process repeats until either the asking price is offered or the queue of buyers is exhausted, in which case the seller can accept the highest offer he has received.

We derive the optimal bidding behavior of buyers and the optimal asking prices set by sellers, characterize the equilibrium, and show that it coincides with the solution to the planner's problem. The fact that our asking price mechanism can implement the efficient allocation is surprising for a number of reasons, as we discuss at the end of Section 4. Chief among them, implementing the planner's allocation requires achieving efficiency along *two* margins: the allocation of the good after buyers arrive, along with the number of buyers that arrive to begin with. However our asking price mechanism affords sellers only a *single* instrument that controls both margins; the asking price determines both the ex-post allocation of the good (by implementing a stopping rule) and the ex-ante expected number of buyers (by specifying the division of the expected surplus).

After establishing that the asking price mechanism implements the solution to the planner's problem, we ask whether sellers would indeed choose to use the asking price mechanism, or if

²As we discuss below, we consider the information made available to buyers to be an endogenous feature of the pricing mechanism.

there exists an alternative mechanism that could potentially increase sellers' profits. To address this question, in Section 5 we consider a similar environment to the one described above, but we allow sellers to select from a more general set of mechanisms. In doing so, we are essentially providing sellers with the option to choose a trading protocol from a large set of extensive form games, in which they are free to make different specifications about the sequence of events, the strategies available to buyers under various contingencies, and even the *information* that buyers have when they select from these strategies.³ In this more general environment, we establish that the asking price mechanism we study in Section 4 maximizes a seller's payoff, regardless of the mechanisms posted by other sellers. As a consequence, there always exists an equilibrium in which all sellers use the asking price mechanism described above. Moreover, while other equilibria can exist, they are all payoff-equivalent; in particular, sellers can do no better than they do in the equilibrium with asking prices. Finally, we show that *any* mechanism that emerges as an equilibrium in this environment will resemble the asking price mechanism along most important dimensions. Therefore, though we cannot rule out potentially complicated mechanisms that also satisfy the equilibrium conditions, the fact that asking prices are both simple and commonly observed suggests that they offer a compelling way to deal with the frictions in our environment.

In Section 6, we flesh out just a few of the model's implications for a variety of observable outcomes. In particular, we study how the asking prices set by sellers, and the corresponding distribution of transaction prices that occur in equilibrium, depend on features of the environment, such as the ratio of buyers to sellers, the degree of ex ante uncertainty in buyers' valuations, and the costs of inspecting the good.

In Section 7, we discuss a few key assumptions and potential extensions of our framework. One particularly noteworthy exercise is showing that a simple variation of our asking price mechanism would produce a distribution of transaction prices that occur below the asking price, a mass point of transactions that occur at the asking price, *and additional transactions that occur above the asking price*. This variation, which preserves all of the normative properties reported above, could be an important step towards understanding, e.g., real estate transactions, which sometimes occur above the asking price. Section 8 concludes, and the Appendix contains all proofs.

Related Literature. We contribute to the literature along two dimensions. The first is normative: we show that our asking price mechanism is both revenue-maximizing and efficient in an

³Indeed, the mechanisms that are available to sellers range from "plain vanilla" pricing schemes—such as auctions, price-posting with rationing, multilateral bargaining games, etc.—to much more complicated games, in which the rules and sequence of play, as well as the information made available to the players, could depend on the number of buyers who arrive at a seller, an individual buyer's place in the queue, the behavior of other buyers, and so on. In short, perhaps the only relevant type of trading mechanism that we do *not* allow are those that depend on the trading mechanisms posted by other sellers, which is standard in this literature; such collusive behavior would seem to violate the competitive spirit that motivates our analysis.

environment with two simple ingredients. The second is positive: we provide a rationale for the use of asking prices and explore the implications for equilibrium prices and allocations. Below, we compare our normative and positive results, respectively, with the existing literature.

As mentioned above, implementing the planner's allocation requires achieving efficiency along two margins. First, a seller's mechanism has to induce socially efficient search behavior by buyers; this requires that each buyer's *ex ante* expected payoff is equal to his marginal contribution to the expected surplus at that seller. Second, a seller's mechanism must maximize the surplus after buyers arrive; that is, the mechanism must ensure that the *ex post* allocation of the good is efficient. These two margins have been studied extensively in separate literatures.

The first margin is a critical object of interest in the literature on competitive (or directed) search, such as Moen (1997), Burdett et al. (2001), Acemoglu and Shimer (1999), and Julien et al. (2000), to name a few. While these papers find that the market equilibrium decentralizes the planner's solution, their results do not extend to our environment. The reason is that, in these papers, the size of the surplus after buyers arrive at a seller is independent of the pricing mechanism. Hence, a simple mechanism (e.g., a single price) can control the division of the expected surplus in an arbitrarily flexible way without any ramifications for *ex post* efficiency. Therefore, within this literature, the papers that are most related to our work are McAfee (1993), Peters and Severinov (1997), Burguet and Sákovics (1999), Eeckhout and Kircher (2010), Virág (2010), Kim and Kircher (2015), Albrecht et al. (2012), and Lester et al. (2015), who consider environments where the pricing mechanism that is posted *does* effect the size of the surplus after buyers arrive. However, in these papers the *ex post* efficient allocation of the good is typically fairly trivial, and hence can be implemented with a simple, simultaneous mechanism (like an auction with no reserve price). In contrast, our model—where buyers incur a cost to learn their valuation—requires that we consider a competitive search model where sellers post *sequential* mechanisms; to the best of our knowledge, this is the first paper to do so.

Sequential mechanisms have received much more attention in the literature that focuses on the second margin highlighted above but abstracts from the first margin, i.e., the literature that studies the allocation of a good when a monopolist seller faces a fixed number of buyers who can learn their valuation at a cost. Within this literature, the paper that is closest to ours is Crémer et al. (2009). Indeed, in our environment, the problem that *each* seller faces after buyers arrive is analyzed as a special case in their paper (Section 3.3). They show that, in this setting, the seller optimally selects a sequential mechanism that requires a series of different take-it-or-leave-it offers and participation fees, both of which may vary according to the number and realizations of previous bids (also see Burguet, 1996).⁴ In a similar environment, Bulow and Klemperer (2009) ignore the

⁴Crémer et al. (2009) also consider richer environments, with various types of heterogeneity. As we discuss below, our asking price mechanism remains optimal after introducing some, but not all types of heterogeneity. For example, if buyers had different inspection costs, as in the general case studied by Crémer et al. (2009), then sellers would want

optimal mechanism and instead focus on two “plain vanilla” mechanisms: a simultaneous auction and a simple sequential mechanism. Choosing between these two mechanisms, they show, involves a trade-off between revenue-maximization and efficiency.⁵

In contrast to these papers, the mechanism we consider is relatively simple, and yet the trade-off in Bulow and Klemperer (2009) does not arise; the asking price is both revenue-maximizing and efficient.⁶ The reason for this difference is twofold. First, it turns out that the information made available to buyers has important effects on the structure of the optimal mechanism. Crémer et al. (2009) and Bulow and Klemperer (2009) assume that prospective buyers observe the behavior of previous bidders who have entered the mechanism, which implies that the optimal mechanism offers different asking prices (and charges different fees) to each buyer who inspects the good. We allow the sellers in our model to control what information is available to buyers, and they choose not to reveal information about previous bidders. This change alone implies that sellers would set the same asking price for all buyers, albeit with an admission fee attached. The second reason our results are different is that the number of buyers who can potentially inspect a seller’s good is exogenous in the papers cited above, whereas in our model it is the endogenous outcome of a game in which sellers compete for buyers. In this competitive setting, as we noted earlier, admission fees are driven to zero in equilibrium.⁷

Turning to the positive results, our explanation is obviously not the only reason why asking prices might be useful.⁸ For one, if buyers are risk averse, an asking price can reduce the uncertainty an individual buyer faces in certain types of auctions, and hence introducing this mechanism can increase a seller’s revenues; see Budish and Takeyama (2001), Mathews (2004), or Reynolds and Wooders (2009). A second explanation for asking prices, which also assumes that buyers incur a cost to learn their valuation, is proposed by Chen and Rosenthal (1996) and Arnold (1999). In these papers, a holdup problem emerges when a buyer and seller bargain over the price after the buyer incurs the inspection cost. An asking price is treated as a ceiling on the bargaining outcome and thus partially solves the holdup problem, providing an ex ante guarantee that the buyer will receive some rents. Our paper is similar in that we also study how a seller’s commitment to a

to use additional devices in order to influence the order in which different types of buyers inspected the good.

⁵See also Schmitz (2003).

⁶Interestingly, the asking price mechanism in our model convexifies, in a way, the binary choice between these two pricing schemes; as we discuss below, one can interpret our mechanism in such a way that the asking price determines the probability of having an auction (after all buyers have inspected the good) as opposed to a sequential sale.

⁷In a different environment, Peters (2001) also finds that sellers who would charge positive admission fees in a monopolistic setting choose not to charge fees in a competitive setting. A similar insight emerges from models that study the problem of a single seller who selects a mechanism, and a large number of buyers who can choose to participate in the mechanism at a cost; see, e.g., Engelbrecht-Wiggans (1987), McAfee and McMillan (1987), and Levin and Smith (1994). The only difference is that the “cost” to a buyer from visiting a seller in a competitive search model is the opportunity cost of not visiting other sellers, which is an equilibrium object, while the cost of participating in these models with one seller is an explicit entry cost, which is exogenous.

⁸Lester (2015) provides a more detailed discussion of existing explanations for asking prices.

pricing mechanism can induce buyers to inspect when they otherwise would not. However, there are a number of important differences, too: in our model, a seller’s pricing decision also affects the number of buyers who arrive to inspect the good and the allocation of the good after inspections occur. Moreover, we show that an asking price is optimal for sellers given almost any mechanism one could imagine, while it’s not clear in Chen and Rosenthal (1996) whether the mechanism they consider is the best way to solve the holdup problem they describe.⁹

A trading mechanism that bears some resemblance to an asking price can also emerge if it is costly for sellers to meet with each buyer, as in McAfee and McMillan (1988). In this case, a seller’s asking price serves as a commitment device to *keep meeting* with buyers until a sufficiently high bid has been received, despite an incentive ex post to stop earlier. This price is playing the opposite role as the asking price in our model; according to our explanation, asking prices serve as a promise to *stop sampling* after a sufficiently high bid, despite an incentive ex post to continue.¹⁰

Finally, asking prices may serve as a device to signal sellers’ private information. For example, in Albrecht et al. (2016), sellers with heterogeneous reservation values use asking prices to signal their type, which allows for endogenous market segmentation. We view this line of research as complementary to our own; certainly the ability of asking prices to signal a seller’s private information, which we ignore, is important.¹¹

2 The Environment

Players. There is a measure θ_b of buyers and a measure θ_s of sellers, and $\Lambda = \theta_b/\theta_s$ denotes the ratio of buyers to sellers. Buyers each have unit demand for a consumption good, and sellers each possess one, indivisible unit of this good. All agents are risk-neutral and ex ante homogeneous.

Search. Buyers can visit a single seller in attempt to trade. As is standard in the literature, frictions arise because buyers must use symmetric strategies. As a result, the number of buyers to

⁹More generally, the various theories described above all consider the problem of a seller in isolation. Therefore, though each one certainly captures a significant component of what asking prices do, they also abstract from something important: the fact that buyers can observe and compare multiple asking prices at once is not only realistic in many markets, but also seems to be a principal consideration when sellers are determining their optimal pricing strategy.

¹⁰To be more precise, McAfee and McMillan (1988) consider a procurement auction in which a buyer meets with a sequence of sellers and each meeting is costly to the buyer; this is equivalent to our environment with the seller incurring the cost of each meeting. Switching which party bears the inspection cost not only changes the entire rationale for an asking price, as described above, but it also changes the nature of the revenue-maximizing mechanism; the solution to our model with this alternative cost structure is available upon request.

¹¹Menzio (2007), Delacroix and Shi (2013), and Kim and Kircher (2015) also study the signalling role of prices in directed search equilibria. However, in papers like Albrecht et al. (2016) and Menzio (2007), asking prices are not uniquely determined, and hence these models are somewhat limited in their ability to draw positive implications about the relationship between asking prices, transaction prices, and market conditions.

arrive at each seller, n , will be distributed according to a Poisson distribution.¹² As is customary in the literature on directed (or competitive) search, we will refer to the expected number of buyers to arrive at a particular seller as the “queue length,” which we denote by λ . As we describe in detail below, the queue length at each seller will be an endogenous variable, determined by the equilibrium behavior of buyers and sellers.

Preferences. All sellers derive utility y from consuming their own good, and this valuation is common knowledge. A buyer’s valuation for any particular good, on the other hand, is not known *ex ante*. Rather, once buyers arrive at a particular seller, they must inspect the seller’s good in order to learn their valuation, which we denote by x . We assume that each buyer’s valuation is an iid draw from a distribution $F(x)$ with support in the interval $[\underline{x}, \bar{x}]$, and that the realization of x is the buyer’s private information. We assume, for simplicity, that $y \in [\underline{x}, \bar{x}]$. This is a fairly weak assumption; the probability that a buyer’s valuation x is smaller than y can be driven to zero without any loss of generality. Moreover, much of the analysis remains similar when $y < \underline{x}$, though the algebra is slightly more involved.

Inspection Costs. After a buyer arrives at a seller, we assume that he must pay a cost k in order to learn his valuation x . Such costs come in many forms, both explicit (i.e., paying for an inspection) and implicit (i.e., the time it takes to learn one’s valuation); we use k to capture all of these costs.¹³

We restrict our attention to the region of the parameter space in which the cost of inspecting the good does not exhaust the expected gains from trade. In particular, we assume that

$$k < \int_y^{\bar{x}} (x - y) f(x) dx. \quad (1)$$

Note that the inequality in (1) does not necessarily imply that a buyer would always *choose* to inspect the good. In what follows, we will assume that a buyer indeed does have incentive to inspect the good before attempting to purchase it, and in Section 7 we derive a sufficient condition to ensure that this is true in equilibrium.

¹²The restriction to symmetric strategies is often motivated by the observation that coordination among agents in a large market seems fairly implausible. Under these restrictions, the number of buyers to arrive at a particular seller follows a binomial distribution when the number of agents is finite, and converges to the Poisson distribution as the number of buyers and sellers tends to infinity. See, e.g., Burdett et al. (2001).

¹³For example, suppose the goods are houses that are roughly equivalent along easily describable dimensions (size, neighborhood, and so on). However, each home has idiosyncratic features that make them more or less attractive to prospective buyers, and these features are only revealed upon inspection; e.g., a family with young children may want to see if the owners of nearby homes also have young children. Quite often, learning one’s true valuation requires more than a quick tour; e.g., an individual who needs to build a home office may hire an architect to estimate how much it will cost. All of these activities are costly, either because they take time or because they require explicit expenditures. Similar costs—of various magnitudes—exist for purchasing a car, renting an apartment, or even hiring an accountant.

Gains from Trade. When n buyers arrive at a seller, the trading protocol in place will determine how many buyers $i \leq n$ will have the opportunity to inspect the good before exchange (potentially) occurs. We normalize the payoff for a buyer who does not inspect, and thus does not trade, to zero. Therefore, if a buyer with valuation x acquires the good after the seller has met with i buyers, the net social surplus from trade is $x - y - ik$. Alternatively, if the seller retains the good for himself after i inspections, the net social surplus is simply $-ik$.

3 The Planner's Problem

In this section, we will characterize the decision rule of a (constrained) benevolent planner whose objective is to maximize net social surplus, subject to the constraints of the physical environment. These constraints include the frictions inherent in the search process, as well as the requirement that buyers' valuations are costly to learn.

The planner's problem can be broken down into two components. First, the planner has to assign queue lengths of buyers to each seller, subject to the constraint that the sum of these queue lengths across all sellers cannot exceed the total measure of available buyers, θ_b . Second, the planner has to specify the trading rules for agents to follow after the number of buyers that arrive at each seller is realized. We discuss these specifications in reverse order.

Optimal Trading Protocol. Suppose n buyers arrive at a seller. As is well-known in the literature (see e.g. Lippman and McCall, 1976; Weitzman, 1979; Morgan and Manning, 1985), it is optimal in this case for the planner to let buyers inspect the good sequentially, and to assign the good immediately whenever a buyer's valuation exceeds a cutoff, x^* , which equates the marginal cost of an additional inspection with the marginal benefit. Importantly, note that the cutoff does not depend on the total number of buyers n , the number of buyers who have inspected the good so far i , or the history of valuations (x_1, \dots, x_i) . The following lemma formalizes this result.

Lemma 1. *Suppose $n \in \mathbb{N}$ buyers arrive at a seller. Letting x^* satisfy*

$$k = \int_{x^*}^{\bar{x}} (x - x^*)f(x)dx, \quad (2)$$

the planner maximizes social surplus by implementing the following trading rule:

- (i) *If $n > 1$ and $i \in \{1, 2, \dots, n-1\}$, the seller should stop meeting with buyers and allocate the good to the agent with valuation $\hat{x}_i \equiv \max\{y, x_1, \dots, x_i\}$ if, and only if, $\hat{x}_i \geq x^*$. Otherwise, the seller should meet with the next buyer.*
- (ii) *If $n=1$ or $i = n$, the seller should allocate the good to the agent with valuation \hat{x}_n .*

Optimal Queue Lengths. We now turn to the optimal assignment of queue lengths across sellers, given the optimal trading protocol once buyers arrive. Notice immediately that the planner's cutoff x^* is not only independent of the number of buyers, n , but also independent of λ , which governs the distribution over n . The reason is that the optimal stopping rule, on the margin, balances the costs and expected gains of additional meetings, conditional on the event that there *are* more buyers in the queue. The probability of this event, per se, is irrelevant: since the seller neither incurs additional costs nor forfeits the right to accept any previous offers if there are no more buyers in the queue, the probability distribution over the number of buyers remaining in the queue—and thus λ —does not affect the planner's choice of x^* .

Let $S(x^p, \lambda)$ denote the expected surplus generated at an individual seller whom the planner assigns a cutoff x^p and a queue length λ . To derive this function, it will be convenient to define

$$q_i(x; \lambda) = \frac{\lambda^i (1 - F(x))^i}{i!} e^{-\lambda(1-F(x))}. \quad (3)$$

In words, $q_i(x; \lambda)$ is the probability that a seller is visited by exactly i buyers, all of whom happen to draw a valuation greater than x (when all buyers' valuations are learned). In what follows, we will suppress the argument λ for notational convenience.

The net surplus generated at a particular seller is equal to the gains from trade, less the inspection costs. To derive the expected gains from trade at a seller with cutoff x^p and queue length λ , first suppose that n buyers arrive at this seller. There are three relevant cases. First, with probability $F(y)^n$ all n buyers draw valuation $x < y$, in which case there are no gains from trade. Second, with probability $F(x^p)^n - F(y)^n$, the maximum valuation among the n buyers is a value $x \in (y, x^p)$, in which case the gains from trade are $x - y$. Note that the conditional distribution of this maximal valuation x has density $[nF(x)^{n-1}f(x)] / [F(x^p)^n - F(y)^n]$. Finally, with probability $1 - F(x^p)^n$, at least one buyer has valuation $x \geq x^p$. In this case, the seller trades with the first buyer he encounters with a valuation that exceeds x^p ; this valuation is a random drawn from the conditional distribution $f(x) / [1 - F(x^p)]$. Taking expectations across values of n , the gains from trade at a seller with cutoff x^p and queue length λ , in the absence of inspection costs, are

$$\begin{aligned} & \sum_{n=1}^{\infty} \frac{e^{-\lambda} \lambda^n}{n!} \left\{ \int_y^{x^p} (x - y) n F(x)^{n-1} f(x) dx + [1 - F(x^p)^n] \int_{x^p}^{\bar{x}} (x - y) \frac{f(x)}{1 - F(x^p)} dx \right\} \\ &= \int_y^{x^p} (x - y) \lambda q_0(x) f(x) dx + [1 - q_0(x^p)] \int_{x^p}^{\bar{x}} (x - y) \frac{f(x)}{1 - F(x^p)} dx. \end{aligned} \quad (4)$$

Now consider the expected inspection costs incurred by buyers at a seller with cutoff x^p and queue length λ . If a buyer arrives at a seller along with n other buyers, he will occupy the $(i + 1)^{th}$ spot in line, for $i \in \{0, \dots, n\}$, with probability $1/(n + 1)$. In this case, he will get to meet with

the seller only when all buyers in spots $1, \dots, i$ draw $x < x^p$, which occurs with probability $F(x^p)^i$. Taking expectations over n implies that the ex ante probability that each buyer gets to meet with a seller with queue length λ , given a planner's cutoff of x^p , is

$$\sum_{n=0}^{\infty} \frac{e^{-\lambda} \lambda^n}{n!} \left\{ \sum_{i=0}^n F(x^p)^i \frac{1}{n+1} \right\} = \frac{1 - q_0(x^p)}{\lambda [1 - F(x^p)]}. \quad (5)$$

Note that this probability approaches one as x^p goes to \bar{x} from below. Moreover, since the expected number of buyers to arrive at this seller is λ , the total *expected* inspection cost incurred by all buyers is simply $([1 - q_0(x^p)] / [1 - F(x^p)])k$.

Using the results above, the expected surplus generated by a seller with queue length λ and stopping rule x^p can be written as

$$S(x^p, \lambda) = \int_y^{x^p} (x - y) \lambda q_0(x) f(x) dx + \frac{1 - q_0(x^p)}{1 - F(x^p)} \left[\int_{x^p}^{\bar{x}} (x - y) f(x) dx - k \right]. \quad (6)$$

Given the optimality of $x^p = x^*$ for any λ , the objective of the planner is then to choose queue lengths at each seller, λ_j for $j \in [0, \theta_s]$, to maximize total surplus $\int_0^{\theta_s} S(x^*, \lambda_j) dj$ subject to the constraint that $\int_0^{\theta_s} \lambda_j dj = \theta_b$. In the following lemma, we establish an important property of the surplus at a particular seller when the optimal trading rule is in place.

Lemma 2. *$S(x^*, \lambda)$ is strictly concave with respect to λ .*

Two factors contribute to the concavity of $S(x^*, \lambda)$ in λ . First, as is standard in models of directed search, the probability that a seller trades is concave in the queue length. This force alone typically implies that the planner assigns equal queue lengths across (homogeneous) sellers. However, in our environment there is an additional force, since the ex post gains from trade are also concave in the number of buyers that arrive: each additional buyer is less likely to meet the seller and, conditional on meeting, is less likely to have a higher valuation than all previous buyers.

The Solution to the Planner's Problem. Taken together, lemmas 1 and 2 are sufficient to establish that the planner maximizes total surplus by assigning equal queue lengths across all sellers, so that $\lambda_j = \Lambda$ for all j . The following proposition summarizes the planner's solution.

Proposition 1. *The unique solution to the planner's problem is to assign equal queue lengths Λ to each seller. After buyers arrive, the planner lets buyers inspect the good sequentially and assigns the good immediately whenever a buyer's valuation exceeds x^* . When all valuations are below x^* and the queue is exhausted, the planner assigns the good to the agent with the highest valuation.*

4 The Decentralized Equilibrium

In this section, we introduce a method of price determination that we call an “asking price mechanism.” We characterize the equilibrium when sellers use this mechanism, establish that this decentralized equilibrium coincides with the solution to the planner’s problem, and discuss our results.

Asking Price Mechanism. An asking price mechanism, or APM, has two components: an asking price and a specific protocol that occurs after an arbitrary number of buyers arrive. Suppose a seller has posted an asking price a and n buyers arrive. The APM dictates that the seller will meet buyers one at a time, with each successive buyer being chosen randomly from the set of remaining buyers, until one buyer bids the asking price or all buyers are met. During a meeting, a buyer incurs the inspection cost k , learns his valuation x , and submits a bid b . Importantly, when a buyer bids, he knows neither the number of other buyers at the seller, $n - 1$, nor his place in the queue.¹⁴

If the buyer bids $b \geq a$, his bid is accepted immediately and trade ensues; the asking price a is the price at which the seller commits to selling his good immediately (and subsequently stops meeting with other buyers). If $b < a$, then the bid is rejected and the seller moves to the next buyer. Finally, if $b \in [y, a)$, then the bid is neither rejected nor immediately accepted. Instead, the seller registers the bid and proceeds to meet the next buyer in line (if there is one). Again, the next buyer incurs the cost k , learns her valuation x , and submits a bid b .

The process described above repeats itself until either the seller receives a bid $b \geq a$, or until he has met with all n buyers. In the latter case, he sells the good to the highest bidder at a price equal to the highest bid, as long as that bid exceeds his own valuation y . A seller who trades at price b receives a payoff equal to b , while a seller who does not trade receives payoff y . The payoff to a buyer who trades at price b is $x - b - k$. A buyer who meets with a seller but does not trade obtains a payoff $-k$. Finally, the payoff of a buyer who does not meet with a seller is equal to zero.

A Buyer’s Strategy and Payoffs. Given an arbitrary distribution of asking prices posted by sellers, we now describe buyers’ optimal search and bidding strategies. We work backwards, first deriving buyers’ optimal behavior and payoffs conditional on meeting a seller who posted an asking price a and attracted a queue λ . Then, given these payoffs, we derive buyers’ optimal search behavior, which determines the queue length λ that corresponds to each asking price a .

To derive buyers’ behavior and payoffs after arriving at a seller, it’s helpful to note that the APM described above is equivalent to a mechanism in which buyers are randomly placed in line; each buyer is sequentially offered the opportunity to purchase the good at the asking price a , in

¹⁴Note that the information available to the buyers should be viewed as a feature of the mechanism. Since we establish below that this mechanism is optimal, it follows that sellers have no incentive ex ante to design a mechanism in which buyers can observe either n or their place in the queue.

which case trade occurs immediately; and, if no buyer chooses to pay the asking price, the good is allocated according to a first-price sealed-bid auction, where buyers are not told how many other buyers are also participating in the auction. As we will see, this alternative interpretation is helpful because it allows us to draw upon well-known results in the auction theory literature.¹⁵

We conjecture, and later confirm, that there exists a cutoff, which we denote x^a , such that buyers with valuation x^a are exactly indifferent between paying the asking price and waiting to (perhaps) participate in an auction. Therefore, if a buyer draws valuation $x \geq x^a$, he will pay a and trade immediately. Otherwise, if $x < x^a$, he will decline and take his chances with the auction.

To derive this cutoff, consider a buyer who has incurred the (sunk) cost k and discovered that his private valuation is x at a seller who has posted an asking price a and has an expected queue length λ , when all other buyers are using a cutoff \tilde{x}^a . To decide whether to accept the asking price or not, the buyer needs to form beliefs about both the probability that the auction will take place and the number of other buyers that will be bidding. Note that when the buyer is asked to inspect, he updates his beliefs regarding both probabilities.¹⁶ In particular, conditional on the buyer meeting the seller (and not paying a), the probability that an auction will take place is

$$\frac{\lambda (1 - F(\tilde{x}^a))}{1 - q_0(\tilde{x}^a)} q_0(\tilde{x}^a). \quad (7)$$

To understand this expression, recall that the probability that a buyer gets to meet a seller is given by (5) and that an auction takes place if, and only if, none of the buyers in the queue draw a valuation above \tilde{x}^a , which happens with (ex ante) probability $q_0(\tilde{x}^a)$. Applying Bayes' rule then yields the expression in (7).

Similar logic can be used to calculate the distribution over the number of competitors that the buyer will face should an auction occur. The total number of buyers at the seller follows a Poisson distribution with mean λ . The auction will take place if all of them have a valuation below \tilde{x}^a , which happens with probability $F(\tilde{x}^a)^n$. Hence, conditional on reaching an auction, the number of competitors equals n with probability

$$e^{-\lambda} \frac{\lambda^n F(\tilde{x}^a)^n}{n! q_0(\tilde{x}^a)} = e^{-\lambda F(\tilde{x}^a)} \frac{(\lambda F(\tilde{x}^a))^n}{n!}. \quad (8)$$

¹⁵This observation is also helpful because it implies that we can replace the first-price auction with any revenue-equivalent auction without changing any of the substantive results below. In most markets, we think the first-price auction most closely resembles the actual method of price determination. However, in Section 7, we explore the implications of using a second-price (ascending bid) auction instead, and argue that this alternative can help rationalize transaction prices above the asking price.

¹⁶In other words, the very act of meeting the seller and inspecting the good is informative, since inspecting the good requires that no previous buyer drew a valuation above \tilde{x}^a . In particular, after inspecting the good, a buyer's posterior belief that many buyers arrived at the seller falls, and he believes it is more likely that he will be competing against relatively few other buyers.

In Lemma 3, below, we use these probabilities to characterize the buyer's optimal bidding strategy and expected payoff should he choose not to pay the asking price, conditional on the asking price a , the queue length λ , and the cutoff \tilde{x}^a being chosen by other buyers. To do so, it will be helpful to define the bidding strategy

$$\widehat{b}(x) = x - \frac{\int_y^x q_0(x') dx'}{q_0(x)}. \quad (9)$$

Lemma 3. *Consider a buyer who arrives at a seller with asking price a and queue length λ , and suppose all other buyers bid $\widehat{b}(x)$ if $x \leq \tilde{x}^a$ and pay the asking price a if $x > \tilde{x}^a$, for some $\tilde{x}^a \in [y, \bar{x}]$. If the buyer discovers a valuation $x \in [y, \bar{x}]$ and chooses not to pay the asking price, the optimal bidding strategy is $\widehat{b}(x)$ for all $x \leq \tilde{x}^a$ and $\widehat{b}(\tilde{x}^a)$ for all $x > \tilde{x}^a$. The ex ante expected payoff from following this strategy is*

$$u(x; \tilde{x}^a) = \begin{cases} \frac{\lambda(1-F(\tilde{x}^a))}{1-q_0(\tilde{x}^a)} \int_y^x q_0(x') dx' & \text{if } x \leq \tilde{x}^a, \\ \frac{\lambda(1-F(\tilde{x}^a))}{1-q_0(\tilde{x}^a)} \left[\int_y^{\tilde{x}^a} q_0(x') dx' + q_0(\tilde{x}^a)(x - \tilde{x}^a) \right] & \text{if } x > \tilde{x}^a. \end{cases} \quad (10)$$

Note that, since $q_0(x)$ is a function of λ , so too are $\widehat{b}(x)$ and $u(x; \tilde{x}^a)$; we have only suppressed λ for notational convenience. Also note that $u(x; \tilde{x}^a)$ is continuous and strictly increasing in x , which confirms our conjecture that buyers follow a threshold strategy.

An equilibrium where buyers play symmetric strategies is thus a threshold x^a such that $u(x^a; x^a) = x^a - a$. From equation (10), then, x^a can be defined by the implicit function

$$a = x^a - \frac{\lambda(1-F(x^a))}{1-q_0(x^a)} \int_y^{x^a} q_0(x) dx. \quad (11)$$

In the Appendix, we confirm that—so long as the asking price is neither too low nor too high—there exists a unique $x^a \in (y, \bar{x})$ that satisfies (11), with $\partial x^a / \partial a > 0$.

Hence, given any a and λ , the buyer's optimal bidding function is completely characterized by

$$b(x) = \begin{cases} 0 & \text{if } x < y \\ \widehat{b}(x) & \text{if } y \leq x < x^a \\ a & \text{if } x^a \leq x, \end{cases} \quad (12)$$

where $\widehat{b}(x)$ is given by (9) and x^a is determined by (11).

Given a buyer's optimal behavior conditional on meeting the seller, we can calculate the ex ante expected utility that a buyer receives from visiting a seller who has posted an asking price a and attracted a queue length λ ,

$$U(a, \lambda) = \frac{1 - q_0(x^a)}{\lambda(1 - F(x^a))} \left[\int_y^{x^a} u(x; x^a) dF(x) + \int_{x^a}^{\bar{x}} (x - a) dF(x) - k \right], \quad (13)$$

where, in a slight abuse of notation, $x^a \equiv x^a(a, \lambda)$ is the implicit function defined in (11).

Given $U(a, \lambda)$, the optimal search behavior of buyers is straightforward: given the set of asking prices that have been posted, along with the search decisions of other buyers, an individual buyer should visit a seller (or mix between sellers) that maximizes $U(a, \lambda)$. More formally, for an arbitrary distribution of posted asking prices, let \bar{U} denote the highest level of utility that buyers can obtain; as is common in this literature, we will refer to \bar{U} as the “market utility.” Then, for any asking price a that has been posted, given the buyers’ optimal threshold and bidding strategy described above, the queue length $\lambda(a) \geq 0$ must satisfy

$$U(a, \lambda(a)) \leq \bar{U}, \text{ with equality if } \lambda(a) > 0. \quad (14)$$

According to equation (14), buyers will adjust their search behavior in such a way to make themselves indifferent between any seller that they visit with positive probability.¹⁷

A Seller’s Strategy and Payoffs. Given the optimal behavior of buyers, we can now characterize the profit-maximizing asking price set by sellers. As a first step, we use the results above to derive the expected revenue of a seller who has set an asking price a and attracted queue length λ :

$$R(a, \lambda) = q_0(y)y + \int_y^{x^a} \hat{b}(x) dq_0(x) + (1 - q_0(x^a))a. \quad (15)$$

Again, note that $x^a \equiv x^a(a, \lambda)$ denotes the optimal cutoff for buyers characterized in (11) and $\hat{b}(x)$ denotes the optimal bidding function characterized in (9). This expression captures the three possible outcomes for a seller: no buyers arrive with valuation $x > y$, in which case the seller consumes his own good; no buyers arrive with valuation $x > x^a$, but at least one buyer has valuation $x > y$, in which case the seller accepts the bid placed by the buyer with the highest valuation; or at least one buyer has a valuation $x \geq x^a$, in which case the seller receives a payoff a .

Sellers want to maximize expected revenue, taking into account that their choice of the asking price a will affect the expected number of buyers that will visit them, λ , through the relationship defined in (14), i.e., $U(a, \lambda) = \bar{U}$. This relationship is akin to a typical demand function: for a given level of market utility, it defines a downward sloping relationship between the asking price a seller sets and the number of customers he receives (in expectation). The seller’s problem can thus be interpreted as a choice over both the asking price and the queue length in order to maximize

¹⁷Following the convention in the literature, if, for some a , there exist multiple values of λ such that $U(a, \lambda) = \bar{U}$, we assume that a seller can post an equilibrium selection device that yields him the greatest profit; see, e.g., McAfee (1993) and Eeckhout and Kircher (2010). Note that the equilibrium we characterize below in Proposition 2 survives even if we relax this assumption; see the proof for additional discussion.

revenue, subject to (14). The corresponding Lagrangian can be written¹⁸

$$\mathcal{L}(a, \lambda, \mu) = R(a, \lambda) + \mu [U(a, \lambda) - \bar{U}]. \quad (16)$$

Equilibrium. In general, an equilibrium is a distribution $G(a, \lambda)$ and a market utility \bar{U} such that (i) every pair in the support of G is a solution to (16), given \bar{U} ; and (ii) aggregating queue lengths across all sellers, given the distribution G and the mass of sellers θ_s , yields the total measure of buyers, θ_b . However, as we establish in the proposition below, in fact there is a unique solution to (16), and hence G is degenerate. Furthermore, this solution coincides with the planner's solution, i.e., the equilibrium is efficient.

Proposition 2. *Given assumption (1), the decentralized equilibrium is characterized by*

$$a = a^* \equiv x^* - \frac{\lambda(1 - F(x^*))}{1 - q_0(x^*)} \int_y^{x^*} q_0(x) dx, \quad (17)$$

$x^a = x^*$, and $\lambda = \Lambda$ at all sellers, with buyers receiving market utility $\bar{U}^* \equiv U(a^*, \Lambda)$. Hence, the decentralized equilibrium coincides with the solution to the planner's problem.

Asking Prices and Constrained Efficiency. Before proceeding, we offer a brief discussion of the result in Proposition 2. Notice immediately that the asking price mechanism we describe has all the right ingredients to implement the efficient allocation at each seller: it elicits buyers' private valuations sequentially, it implements a stationary cutoff rule, and it allows for the seller to trade with any buyer that has inspected the good (i.e., it allows for perfect recall).

What is less clear is whether the asking price that sellers choose in equilibrium will implement the *efficient* cutoff, x^* . The reason is that the choice of a is affecting two margins: the expected number of buyers who arrive at the seller, λ , and the stopping rule after buyers arrive, x^a . One way to understand these two margins more clearly is to use the relationship

$$R(a, \lambda) = S(x^a(a, \lambda), \lambda) - \lambda U(a, \lambda)$$

to rewrite the Lagrangian (16), where $x^a(a, \lambda)$ is defined in (11). Taking first order conditions and substituting the constraint, (14), we get that profit maximization requires

$$\frac{\partial \lambda}{\partial a} \left[\frac{\partial S}{\partial \lambda} - \bar{U} \right] + \frac{\partial S}{\partial x^a} \frac{dx^a}{da} = 0. \quad (18)$$

¹⁸So long as \bar{U} is sufficiently small, the solution will be interior and hence the first-order conditions of the Lagrangian are necessary. Sufficiency follows from the fact that the unique solution coincides with the planner's allocation, as we prove below.

The first term in (18) represents the seller’s marginal revenue from increasing a , holding the threshold x^a constant. In particular, note that the marginal revenue from an additional buyer, in expectation, is equal to the additional surplus this buyer creates, $\partial S/\partial\lambda$, less the “cost” of acquiring this buyer, \bar{U} . The second term represents the seller’s marginal revenue from increasing a , holding the queue length λ constant. Profit maximization, of course, just requires that the *sum* of the two terms in (18) is equal to zero. Constrained efficiency, on the other hand, requires that *each* of the two terms in (18) is equal to zero: implementing the solution to the planner’s problem requires both maximizing the surplus after buyers arrive and inducing the right number of buyers to arrive in the first place. The former requirement demands that $\partial S/\partial x^a = 0$, so that the asking price implements the stopping rule $x^a = x^*$. The latter demands that, in equilibrium, $\partial S/\partial\lambda = \bar{U}$, which is a standard condition for the ex ante efficient allocation of buyers across sellers: it implies that each buyer receive (in expectation) his marginal contribution to the match.¹⁹

A priori, there is no reason to think that one value of a will ensure that both of these conditions are satisfied. Indeed, since one instrument is affecting two margins, one could easily imagine a seller being forced to trade off ex post efficiency—so that $\partial S/\partial x^a \neq 0$ —in order to influence ex ante demand.²⁰ And yet, this doesn’t happen. Instead, the asking price that implements the efficient stopping rule also delivers to each buyer his expected marginal contribution to the match surplus.

Before explaining why, it’s helpful to note that a similar result emerges in an environment with no inspection costs, i.e., with $k = 0$. In this case, the efficient allocation can be implemented with a standard (first- or second-price) auction with a reserve price equal to y , which also happens to yield each buyer his expected marginal contribution to the match surplus.²¹

In our environment, with $k > 0$, the connection between the asking price mechanism, the efficiency of the stopping rule, and the buyers’ ex ante expected payoffs is much more complicated. To understand why the asking price that implements x^* also yields each buyer his expected marginal contribution to the match surplus, it’s helpful to consider the possible valuations realized by a marginal buyer.

First, suppose the buyer has a valuation $x < x^*$. In this case, the buyer only contributes to the match surplus if his valuation is the maximum among all buyers, in which case the marginal contribution is the difference between his valuation and the next highest valuation. Since the good is allocated through an auction when the maximum valuation is less than x^* , the logic above (when $k = 0$) implies that this buyer’s payoffs coincide with his marginal contribution to the surplus.

¹⁹See, e.g., Mortensen (1982), Hosios (1990), and Moen (1997).

²⁰If the seller had two instruments at his disposal—say, an asking price to implement a stopping rule and a fee (or subsidy) to transfer rents at will—clearly the seller would choose an asking price to implement x^* and a fee or subsidy to independently influence the queue length λ . What’s interesting here is that such fees or subsidies are not necessary.

²¹This is a well known result in auction theory. To convey the intuition most clearly, it’s easiest to consider a second price auction, where a buyer is rewarded the good when he has the maximal valuation and his payoff is exactly the difference between his own valuation and the maximum valuation that would be realized in his absence. By standard revenue equivalence results, the same is true *in expectation* in a first-price auction.

Second, suppose the buyer has a valuation $x = x^*$. By accepting the asking price a^* , this buyer affects the social surplus in two different ways. First, in the event that no other buyer has a valuation greater than x^* , this buyer increases the maximum valuation from $\mathbb{E}[\hat{x}|\lambda, \hat{x} < x^*]$ to x^* , where \hat{x} denotes the maximum valuation among n buyers and n is drawn from a Poisson distribution with mean λ . Second, by accepting the asking price, this buyer prevents other buyers from inspecting the good, which has both costs and benefits: it saves some buyers from incurring the cost k and discovering a low valuation $x < x^*$, but prevents other buyers from discovering a high valuation $x > x^*$ and potentially acquiring the good. However, recall that x^* is exactly the threshold that offsets these costs and benefits, leaving only the first effect on the social surplus. This effect is easily identifiable in the payoff of a buyer with valuation x^* , which can be written

$$x^* - a^* = \frac{\lambda[1 - F(x^*)]q_0(x^*)}{1 - q_0(x^*)} \int_y^{x^*} \frac{q_0(x')}{q_0(x^*)} dx'. \quad (19)$$

The first factor on the right-hand side is the probability that no other buyer has valuation $x \geq x^*$, and the second factor is equal to $x^* - \mathbb{E}[\hat{x}|\lambda, \hat{x} < x^*]$. Hence, this buyer's payoffs coincide exactly with his marginal contribution to the surplus.

Similar logic holds when the buyer's valuation is $x > x^*$. Again, since the threshold is x^* , the costs and benefits of preventing later buyers from inspecting the good exactly offset, leaving only the effect of this buyer increasing the maximum valuation from $\mathbb{E}[\hat{x}|\lambda, \hat{x} < x^*]$ to x . Since accepting the asking price yields a payoff $x - a^* = (x^* - a^*) + (x - x^*)$, this buyer receives the share of the surplus that a buyer with valuation x^* receives, $x^* - a^*$, *plus* the additional surplus he creates, $x - x^*$. Hence, again, this buyer's payoffs coincide exactly with his marginal contribution to the surplus.

5 General Mechanisms

Proposition 2 characterizes the equilibrium that arises when sellers compete by posting asking prices, and establishes that the equilibrium coincides with the solution to the social planner's problem. However, it remains to be shown that sellers would in fact *choose* to utilize the asking price mechanism if we expanded their choice set to include more general mechanisms.

In this section, we establish several key results. First, even when sellers are free to post arbitrary mechanisms, all sellers posting an asking price mechanism with $a = a^*$ (described in Proposition 2) remains an equilibrium. Moreover, while other equilibria can arise, all of these equilibria are payoff-equivalent to the equilibrium with optimal asking prices; in particular, there is no equilibrium in which sellers earn higher payoffs than they do in the equilibrium with optimal asking prices. Finally, *any* mechanism that emerges as an equilibrium in this environment will resemble

the asking price mechanism along several important dimensions: any equilibrium mechanism will require that the seller meets with buyers sequentially, that meetings continue until a buyer draws a valuation x^* , and that the (expected) payment by a buyer with valuation $x \geq x^*$ is equal to a^* . Hence, though we cannot rule out potentially complicated mechanisms that satisfy these properties, the fact that asking prices are both simple and commonly observed suggests that they offer a compelling way to deal with the frictions in our environment.

Mechanisms. Consider a seller who receives n buyers. In general, a mechanism is going to specify an extensive form game that determines, for each $i \in \{1, \dots, n\}$,

1. a probability $\phi_{i,n} : \mathcal{X}^{i-1} \rightarrow [0, 1]$ that the i^{th} buyer inspects the good conditional on the messages $(\tilde{x}_1, \dots, \tilde{x}_{i-1}) \in \mathcal{X}^{i-1}$ reported by previous buyers, where $\mathcal{X} = \{\emptyset\} \cup [\underline{x}, \bar{x}]$ is the space of messages available to buyers, $\tilde{x} = \emptyset \in \mathcal{X}$ denotes the event that buyer i did not inspect the good, and $\mathcal{X}^0 = \{\emptyset\}$ by convention;
2. a disclosure rule $\sigma_{i,n} : \mathcal{X}^{i-1} \rightarrow \Sigma$ that specifies the signal that the i^{th} buyer receives conditional on the vector of reports from previous buyers, $(\tilde{x}_1, \dots, \tilde{x}_{i-1}) \in \mathcal{X}^{i-1}$, where Σ is the space of signals that can be disclosed;²²
3. a decision rule $\delta : \Sigma \rightarrow [0, 1]$ that specifies the probability that a buyer will inspect the good conditional on receiving the signal $\sigma \in \Sigma$;
4. a report $\tilde{x} : \Sigma \times \mathcal{X} \rightarrow \mathcal{X}$ that the buyer sends to the seller conditional on receiving a signal $\sigma \in \Sigma$ and either not inspecting or discovering a valuation $x \in [\underline{x}, \bar{x}]$;

along with

5. an allocation rule $\alpha_n : \mathcal{X}^n \rightarrow \{0, 1, \dots, n\}$ which specifies that the good is allocated to either the i^{th} buyer in the queue, for $i \in \{1, \dots, n\}$, or the seller ($i = 0$), conditional on the set of messages $(\tilde{x}_1, \dots, \tilde{x}_n)$;
6. transfers $\tau_n : \mathcal{X}^n \rightarrow \mathbb{R}^n$ which specifies the transfers to/from each of the n buyers, conditional on the set of messages $(\tilde{x}_1, \dots, \tilde{x}_n)$;

subject to standard feasibility and aggregate consistency constraints.

We will restrict attention to the set of all mechanisms \mathcal{M} that satisfy individual rationality and incentive compatibility. In particular, for every $m \in \mathcal{M}$, a buyer's expected payoff when instructed to inspect the good is non-negative, given the information available to him, and truthfully reporting

²²The set of potential signals is quite vast. For one, the seller could disclose the buyer's exact place in line, some information about his place in line (i.e., whether he is among the first n buyers) or no information about his place in line. Similarly, the seller could disclose all of the valuations reported by previous buyers, no information about previous buyers, or various statistics summarizing the reports that have been received, such as the maximum valuation reported or whether any buyers had reported a valuation above some threshold.

his valuation x (at least weakly) dominates reporting any $x' \neq x$.²³ In very similar environments, Crémer et al. (2009) and Pans (2013) show that these two additional restrictions—namely, that buyers are both *obedient* when instructed to inspect and *truthful* in reporting their type—are essentially without loss of generality.²⁴ Given these implicit restrictions, a mechanism $m \in \mathcal{M}$ can simply be summarized by the inspection probabilities and rules for disclosure, allocations, and transfers described above.

Payoffs. Let the ex ante expected payoff of a seller who posts a mechanism $m \in \mathcal{M}$ and attracts a queue λ be denoted by $\mathbf{R}(m, \lambda)$, and let the expected payoff of each buyer in his queue be denoted by $\mathbf{U}(m, \lambda)$. The total payoffs (net of y) cannot exceed the amount of surplus $\mathbf{S}(m, \lambda)$ generated by the seller’s chosen mechanism, which in turn cannot exceed the surplus created by the mechanism that implements the planner’s solution, $\mathbf{S}^*(\lambda) \equiv S(x^*, \lambda)$. That is, for all m and λ ,

$$\mathbf{R}(m, \lambda) + \lambda \mathbf{U}(m, \lambda) - y \leq \mathbf{S}(m, \lambda) \leq \mathbf{S}^*(\lambda). \quad (20)$$

Clearly, Pareto optimality requires that the surplus be divided between the seller and the buyers, so we restrict attention to mechanisms that satisfy the first condition in (20) with equality. We call a mechanism that satisfies the second condition in (20) with equality *surplus-maximizing*.

Equilibrium. An equilibrium in this more general environment is a distribution of mechanisms $m \in \mathcal{M}$ and queue lengths $\lambda \in \mathbb{R}_+$ across sellers, along with a market utility \bar{U} , such that (i) given \bar{U} , each pair (m, λ) maximizes profits $\mathbf{R}(m, \lambda)$ subject to the constraint $\mathbf{U}(m, \lambda) = \bar{U}$; and (ii) aggregating queue lengths across all sellers yields the total measure of buyers, θ_b . Given this definition, we now establish that a mechanism $m \in \mathcal{M}$ is an optimal strategy if, and only if, it creates the same surplus and the same expected payoffs as the asking price mechanism described in the previous section, where the asking price a is set to implement the cutoff x^* .

Proposition 3. *Take any candidate equilibrium with market utility $0 < \bar{U} < \int_y^{\bar{x}} (x - y) f(x) dx - k$. Let m_a^* denote the asking price mechanism that implements the cutoff x^* , and let λ_a satisfy $\mathbf{U}(m_a^*, \lambda_a) = \bar{U}$. A mechanism $m \in \mathcal{M}$, which attracts a queue length λ , maximizes a seller’s expected profits if, and only if, $\mathbf{S}(m, \lambda) = \mathbf{S}^*(\lambda)$, $\mathbf{R}(m, \lambda) = \mathbf{R}(m_a^*, \lambda_a)$, and $\lambda = \lambda_a$.*

The intuition behind Proposition 3 is illustrated in Figure 1. Consider a candidate equilibrium with market utility \bar{U} in which a seller posts a mechanism m_1 and receives a queue λ_1 satisfying $\bar{U} = \mathbf{U}(m_1, \lambda_1)$, yielding expected profits $\mathbf{R}(m_1, \lambda_1) - y = \mathbf{S}(m_1, \lambda_1) - \lambda_1 \bar{U}$.

²³We are also implicitly assuming *truthful disclosure*, i.e., the seller does not report false information to the buyers.

²⁴There are, however, two implicit restrictions on the mechanism space which are worth noting, though both are standard in the literature on competing mechanisms (see, e.g., Eeckhout and Kircher, 2010). First, we do not allow mechanisms to condition on buyers’ ex ante identities; identical buyers must be treated symmetrically. Second, we do not allow mechanisms to condition on the trading mechanisms posted by other sellers.

INSERT FIGURE 1 HERE

The first result is that m_1 must be surplus-maximizing. To see why, suppose $\mathbf{S}(m_1, \lambda_1) < \mathbf{S}^*(\lambda_1)$, corresponding to point 1 in Figure 1. This seller’s expected profits correspond to the intersection of the vertical axis with the line through point 1 with slope \bar{U} . Clearly this cannot be consistent with equilibrium behavior: the seller could deviate to a surplus-maximizing mechanism m_2 that attracts the same queue length λ_1 but yields a larger surplus $\mathbf{S}(m_2, \lambda_1) = \mathbf{S}^*(\lambda_1)$, corresponding to point 2.²⁵ Since this deviation increases the size of the surplus, while holding constant the buyers’ market utility, it strictly increases the seller’s profits, i.e., $\mathbf{R}(m_2, \lambda_1) > \mathbf{R}(m_1, \lambda_1)$.

In general, a seller can obtain an even higher payoff. As the figure shows, profits are maximized at point 3. That is, not only must any equilibrium mechanism lie on the surplus-maximizing frontier $\mathbf{S}^*(\lambda)$, it must also induce a queue length λ_3 such that²⁶

$$\left. \frac{d\mathbf{S}^*(\lambda)}{d\lambda} \right|_{\lambda=\lambda_3} = \bar{U}. \quad (21)$$

Equation (21) is a typical requirement for profit-maximizing behavior: the left-hand side is the marginal benefit of attracting a longer queue length, while the right-hand side is the marginal cost.

Importantly, the asking price mechanism that implements x^* satisfies this condition, as we have demonstrated in the previous section. Hence, for any market utility \bar{U} , a seller (weakly) maximizes his revenue if he posts a surplus-maximizing asking price mechanism. In other words, irrespective of the behavior of other sellers, it is always optimal for an individual seller to post an asking price mechanism that implements x^* . Note that—out of equilibrium—the asking price that achieves this may not be equal to a^* , since that requires $\bar{U} = \bar{U}^*$.

Finally, any profit-maximizing mechanism must attract a queue length equal to the queue length at a seller who posts m_a^* . Intuitively, since the queue length λ_a equates the marginal benefit of attracting more buyers with the more marginal cost, any other queue length would yield the seller strictly lower profits. Therefore, in equilibrium, all sellers must attract the same queue length, so that $\lambda = \Lambda$. The following corollary is an immediate consequence of the results in Proposition 3.

Corollary 1. *All sellers posting the optimal asking price mechanism m_a^* and attracting a queue Λ is an equilibrium within the mechanism space \mathcal{M} .*

Therefore, even when sellers are free to post arbitrary mechanisms, posting an asking price mechanism with $a = a^*$ is consistent with equilibrium behavior. Now, it is true that other mechanisms could also be utilized in equilibrium, but Proposition 3 implies that any such mechanism

²⁵Such a deviation could be achieved by setting an asking price that implements the cutoff x^* , along with a fee (or subsidy) that ensures the expected payoff to buyers—and hence λ_1 —was unchanged.

²⁶Recall that $\mathbf{S}^*(\lambda) = S(x^*, \lambda)$ is strictly concave in λ .

will be similar to the asking price mechanism along several important dimensions. To start, since any equilibrium mechanism must be surplus-maximizing, and thus implement an allocation that coincides with the unique solution to the planner’s problem, the mechanism must feature sequential meetings between the seller and the buyers, with a stopping rule x^* . Moreover, since the ex ante probability that each buyer gets to meet with the seller must be equal in any equilibrium, and any mechanism that arises in equilibrium must also be payoff-equivalent to the equilibrium with optimal asking prices, it follows that the expected payment by buyers with valuation $x \in [x^*, \bar{x}]$ must equal the optimal asking price. Therefore, even when sellers utilize an alternative mechanism in equilibrium, the expected transfer from a buyer who “stops” the sequential inspection process will, indeed, equal a^* . The following corollary summarizes.

Corollary 2. *Any equilibrium strategy $m \in \mathcal{M}$ must specify that the seller meet with buyers sequentially; that these meetings stop if, and only if, either the buyer draws a valuation $x \geq x^*$ or the end of the queue is reached; and that the expected payment for a buyer who draws valuation $x \geq x^*$ is equal to a^* .*

6 Comparative Statics

In this section, we study the equilibrium allocation and distribution of prices, and analyze how they change with features of the economic environment, such as the ratio of buyers to sellers, the cost of inspecting the good, and the degree of ex ante uncertainty in buyers’ valuations.

Figure 2 below plots a typical CDF of transaction prices, where $b = 0$ represents sellers that do not trade.²⁷ Notice that a fraction $q_0(y) = e^{-\Lambda[1-F(y)]}$ of sellers do not trade, either because no buyers arrive, which occurs with probability $q_0(\underline{x}) = e^{-\Lambda}$, or because $n \geq 1$ buyers arrive but their valuations do not exceed the seller’s valuation y , which occurs with probability $q_0(y) - q_0(\underline{x})$. A fraction $q_0(x^*) - q_0(y)$ of sellers ultimately accept a bid b that is strictly less than the asking price. Letting $\hat{x}(b) = \hat{b}^{-1}(x)$, where $\hat{b}(x)$ is defined in (9), the (cumulative) distribution of winning bids $b \in [y, \hat{b}(x^*)]$ is simply $q_0(\hat{x}(b))$. Finally, a fraction $1 - q_0(x^*)$ of sellers trade at the asking price.

INSERT FIGURE 2 HERE

Notice that there is a mass point of transactions that occur at a^* and a gap in the distribution between $\hat{b}(x^*)$ and a^* . Intuitively, it cannot be optimal for a buyer to offer a price arbitrarily close to the asking price; such a strategy would be dominated by offering $b = a^*$, which implies a discrete increase in the probability of trade at the cost of an arbitrarily small increase in the terms of trade.

²⁷For the sake of illustration, all numerical examples below have been generated with the assumption that x is uniformly distributed, though all of the results are true for arbitrary distributions.

Figures 3–5 illustrate how equilibrium prices are affected by changes in the ratio of buyers to sellers (Λ), changes in the inspection cost (k), and changes in the amount of ex ante uncertainty about a buyer’s valuation (the support of the F), respectively.

Figure 3 illustrates that an increase in Λ causes a decrease in the fraction of sellers who do not trade and an increase in the fraction of sellers who trade at the asking price. Moreover, an increase in Λ causes an increase in the asking price, and a first-order stochastic dominant shift in the distribution of transaction prices. Hence, as in standard models of competitive search, an increase in the buyer-seller ratio leads to higher prices in equilibrium. Notice, however, that the degree of dispersion in prices will be non-monotonic in Λ : though price dispersion exists for intermediate values of Λ , the equilibrium price distribution becomes degenerate at $b = y$ ($b = x^*$) as Λ converges to zero (infinity).

Figure 4 illustrates that a decrease in the inspection cost k causes both the asking price a^* and the cut-off x^* to increase. However, the buyers’ bidding function $\widehat{b}(x)$ is unaffected, and hence the lower tail of the equilibrium price distribution is unaffected. As a result, a decrease in k leads to fewer transactions at the asking price. Finally, as k converges to zero, the optimal x^* converges to \bar{x} and the pricing mechanism converges to a standard first-price auction.

Lastly, Figure 5 illustrates the effect of decreasing ex ante uncertainty about a buyer’s valuation: the figure plots the equilibrium distribution of prices when the distribution of valuations $F(x)$, with support $[\underline{x}, \bar{x}]$, is replaced with the truncated distribution F' with support $[\underline{x}', \bar{x}]$, where $\underline{x}' > \underline{x}$.²⁸ Notice that an increase in the lower bound of the support of F implies that fewer sellers don’t trade, even though equilibrium queue lengths remain $\lambda = \Lambda$ at each seller, since it is less likely for a seller to be visited by buyers who have a valuation strictly less than y . Also notice that prices increase for several reasons. For one, sellers set higher asking prices; on the margin, the expected gain from meeting with an additional buyer is larger since the truncated distribution F' first-order stochastically dominates the original distribution F . Moreover, since *other* buyers are more likely to draw a high valuation, there is more competition amongst buyers. This puts upward pressure on the bidding function, further increasing transaction prices. Given these two forces, clearly sellers’ profits increase. Buyers, on the other hand, are more likely to trade, but at less favorable prices.

INSERT FIGURES 3–5 HERE.

²⁸One interpretation of this exercise is that a new technology (e.g., the Internet) replaces an old technology (e.g., the newspaper), allowing buyers to learn more information about each seller’s good before choosing a seller to visit. Given such a technology, buyers could avoid visiting sellers where they were sure to draw a low valuation (e.g., $x \in [\underline{x}, \underline{x}']$), and only visit sellers for which $x \in [\underline{x}', \bar{x}]$.

7 Assumptions and Extensions

In this section, we discuss several of our key assumptions, along with a few potentially interesting extensions of our basic framework.

Transactions Above the Asking Price. In the equilibrium characterized in section 4, no transaction takes place at a price that exceeds the asking price. While this prediction rings true in some markets, it is perhaps less desirable for the analysis of other markets, such as housing, where transactions sometimes occur above the asking price. It is therefore important to emphasize that our model can easily explain such transaction prices.²⁹ Recall that one can interpret our asking price mechanism as a two-stage process in which the seller first sequentially offers the good to each buyer at a price a^* and then organizes a first-price auction if no buyer accepts this offer. Given Proposition 3 and standard revenue-equivalence results, an alternative, optimal mechanism would be to replace the first-price auction with a second-price auction in the second stage of the game.

Note that such a mechanism has a natural interpretation in the context of a housing market. It corresponds to a scenario in which a seller, after getting a number of offers which are all below the asking price, contacts the buyers again, informs them of the competing bids, and asks them whether they would like to increase their offer. In that case, a “bidding war” ensues until a single buyer remains, implementing the outcome of a second-price auction.

INSERT FIGURE 6 HERE

While this mechanism yields the same *expected* payoffs as our (first-price) asking price mechanism, the distributions of *realized* transaction prices will differ. Figure 6 shows a typical CDF of transaction prices, where again we use a price equal to 0 to represent sellers that do not trade. As buyers bid their valuation in a second-price auction, a transaction price above the asking price arises when at least two buyers have a valuation above a but below x^a (or otherwise one of them would end the game by paying the asking price in the first stage). As we report in Lemma 4 below, the probability of this event is strictly positive because $a < x^a$. Hence, this modification yields a micro-founded model in which a positive mass of transactions occur at the posted asking price, while other transactions occur both below and above the asking price.

Lemma 4. *Suppose all sellers post a two-stage mechanism consisting of an asking price in the first stage, followed by a second-price auction. The probability that a transaction takes place at a price exceeding the asking price then equals*

$$e^{-\lambda(1-F(x^a))} - e^{-\lambda(1-F(a))} - \lambda e^{-\lambda(1-F(a))} (F(x^a) - F(a)) > 0.$$

²⁹We would like to thank Rob Shimer for pointing out this feature of our model.

Commitment. Throughout our analysis, we also assume that sellers can commit to carrying out the mechanism that they post. In particular, when a seller posts an asking price, we assume that he commits to trading with the first buyer who offers to pay this price, even though *ex post* he would prefer to renege and meet with all buyers.

Though this assumption is strong, we believe there are a number of ways to enforce this type of behavior. Some are technological. For example, online auction sites like eBay and Amazon allow sellers to pre-commit to an asking price (what they call “Buy-It-Now” or “Take-It” prices, respectively) in which the auction immediately stops once this price offer is received. In other cases, there are institutions that make it costly to renege on an asking price. For example, as Stacey (2012) points out, real estate agents can serve as commitment devices in the U.S. housing market, since sellers are typically required to pay their agent’s commission if they receive a *bona fide* offer at the asking price, whether or not the offer is accepted.³⁰ Finally, even without such technologies or institutions, it’s well known that reputation can sustain commitment in markets with repeated interactions. For all of these reasons, we think that the assumption of full commitment is a reasonable approximation of the way goods are sold in certain markets.

Endogenous Inspection. Throughout the text we assumed that buyers inspect the good and learn their valuation before submitting a bid. One interpretation of this assumption is that it is a technological constraint: a buyer simply must go and meet with the seller in order to make an offer (say, he needs to sign certain documents), and this process is costly. However, for many applications it may be more appropriate to treat the decision to inspect the good as endogenous. In such an environment, if the inspection cost k is too large, there may be circumstances under which the buyer (or the planner) prefers to forgo inspection and place a bid (or trade) without knowing the valuation. The following lemma derives a condition on k to ensure that this is never the case.

Lemma 5. *The planner always instructs buyers to inspect the good upon meeting a seller and, in the decentralized equilibrium, buyers always choose to inspect the good before submitting a bid if and only if*

$$k < \int_x^y (y - x) f(x) dx. \tag{22}$$

Hence, the analysis in Sections 3 and 4 is consistent with an environment in which the decision to inspect the good is endogenous, but k is sufficiently small to satisfy (22). In words, this inequality implies that inspection is always optimal if the cost of inspecting is smaller than the costs associated with inefficient trade, which occurs when the seller values the good more than the buyer who receives it. For many goods that are sold with asking prices—such as houses or cars—the

³⁰Indeed, in other countries—such as France—a seller is legally required to accept a full price offer. We thank an anonymous referee for pointing this out.

assumption that the inspection cost k is small relative to the potential gains (or losses) from trade seems to be appropriate. However, it's worth noting that even when (22) is violated, buyers will still almost always inspect the good before bidding in large regions of the parameter space.

Information Structure. In the asking price mechanism we consider, sellers choose not to reveal a buyer's place in line or the number of other buyers that have arrived. This turns out to be exactly the right amount of information to reveal to the buyer in order to ensure that such a simple, stationary mechanism can implement the optimal stopping rule.

To see this, suppose we replace our asking price mechanism with a standard (sealed-bid) auction. In this case, since buyers bid simultaneously, they have *no information* about the valuations of other buyers, and one can show that too many buyers ultimately inspect the good. Now suppose that we allow each buyer to observe the bids of those who have inspected the good before him, as in Bulow and Klemperer (2009). Then, in the simple sequential mechanism they study, each buyer has an incentive to place a bid that inefficiently preempts further participation. Therefore, if buyers have *too much information*, then too few buyers will ultimately inspect the good.³¹

With the equilibrium asking price, a buyer who is about to inspect the good is able to infer *just enough*: namely, that no other buyer has had a valuation high enough to bid the asking price, making it worthwhile to pay the inspection cost.³² However, he is not able to infer more, which eliminates the strategic motive to bid higher in attempt to dissuade future entrants. In this sense, the equilibrium asking price is essentially a “jump-bid” that does pre-empt entry, but precisely when it is efficient to do so.

Heterogeneity. Before concluding, we discuss the robustness of our results to the introduction of various types of heterogeneity. To start, if the goods for sale are heterogeneous, and this heterogeneity is observable *ex ante*, then our results are essentially unchanged.³³ To be more precise, suppose there are I types of goods. It could be that buyers value each type of good differently: in the context of our model, buyers could draw a valuation x from $F_i(x)$ for each $i \in \{1, \dots, I\}$. It could be that sellers value each type of good differently: in the context of our model, sellers could have a valuation y_i for a good of type $i \in \{1, \dots, I\}$. Or it could be both. As long as the differences across goods are verifiable *ex ante*, then our asking price mechanism is still an equilibrium, and the equilibrium is efficient.

If buyers are *ex ante* heterogeneous, however, the optimal mechanism would likely have more

³¹This is precisely why Crémer et al. (2009) find that the optimal mechanism must include asking prices and fees that vary with each buyer in the queue, which can be used to correct the distortions in the simple mechanism of Bulow and Klemperer (2009).

³²In a related environment with two bidders, Pans (2013) also finds that partial disclosure is optimal for the seller. McAdams (2015) also studies optimal disclosure in an environment where buyers choose when to enter a mechanism.

³³For a formal proof, see the working paper version Lester et al. (2013).

moving parts. For example, if buyers had different inspection costs—say, $k_1 \neq k_2$ —then the optimal mechanism would have two asking prices and a fee. This would ensure that buyers with lower inspection costs could inspect first, in exchange for a fee, and then buyers with higher inspection costs would inspect second (so long as no buyer in the first group paid the initial asking price). Within each group of buyers, however, a single asking price would play exactly the same role that it plays in our model.

Finally, if sellers (or their goods) are *ex ante* heterogeneous and this is *unobservable ex ante*, then the mechanisms that are posted—including an asking price—can potentially signal information about the type of the seller or his good. We view this signalling role as quite distinct, though complementary, to the role that asking prices play in our model.

8 Conclusion

In most economic models, it is assumed that either (i) non-negotiable prices are set by sellers; (ii) prices are the outcome of a bargaining game; or (iii) prices are determined through an auction. However, many goods (and services) are sold in a manner that is not consistent with any of these three pricing mechanisms, but rather seems to combine elements of each. First, sellers announce a price, as in models with price-posting, at which they're willing to sell their good or service immediately. However, as in bargaining theory, buyers can submit a counteroffer that will also be considered by the seller. Finally, in the event that no buyer offers the asking price within a certain period of time, the set of counteroffers that the seller has received are aggregated, as in an auction, and the object is awarded to the buyer with the best counteroffer. Despite the prevalence of this pricing scheme in many markets, it has received little attention in the academic literature. In this paper, our objective was to construct a sensible economic model to help us understand how and why this type of pricing mechanism can be an efficient way of selling goods and services. Combining two simple, realistic ingredients—namely, competition and costly inspection—we showed that asking prices emerge as the mechanism that is both revenue-maximizing and efficient. As a result, the framework developed here provides a theory of asking prices that is both micro-founded and tractable, offers a variety of testable predictions, and lends itself easily to various extensions.

Appendix A: Proofs

Proof of Lemma 1

Proving Lemma 1 requires establishing optimality of both the sequential search strategy and the stopping rule with constant cutoff x^* . Both are well-known results in the literature, so we sketch the intuition below and refer the reader to e.g. Lippman and McCall (1976), Weitzman (1979) and Morgan and Manning (1985) for a more rigorous treatment.

Optimality of Sequential Search. To see that it can never be optimal to learn the valuation of more than one buyer at a time, suppose $n \geq 2$ buyers arrive at a seller, and consider the planner's decision of whether to learn the valuations of the first two buyers sequentially or simultaneously. Let $Z_i(\hat{x}_i) - y$ denote the net expected surplus from continuing to learn buyers' valuations (under the optimal policy) given that the maximum valuation of the seller and the i buyers sampled so far is $\hat{x}_i \equiv \max\{y, x_1, \dots, x_i\}$.

The net social surplus from learning the first two buyers' valuations simultaneously is then

$$-2k - y + \max\{\hat{x}_2, Z_2(\hat{x}_2)\}, \quad (23)$$

whereas the surplus from learning these valuations sequentially is

$$-k - y + \max\{\hat{x}_1, -k + \max\{\hat{x}_2, Z_2(\hat{x}_2)\}\} = -2k - y + \max\{\hat{x}_1 + k, \hat{x}_2, Z_2(\hat{x}_2)\}. \quad (24)$$

Clearly, the expression in (24) is weakly larger than (23) for any x_1 and x_2 , and strictly larger for some x_1 and x_2 . Taking the expectation over all possible realizations therefore implies that a planner would want to learn these valuations sequentially ex ante. It is straightforward to extend this argument to learning the valuations of $j > 2$ buyers simultaneously after any number $i \leq n - j$ buyers have already inspected the good.

Optimality of the Stopping Rule. Suppose the seller is visited by $n \in \mathbb{N}$ buyers. Let $Z_i(\hat{x}_i)$ denote the expected surplus from sampling the $(i+1)^{th}$ buyer, for $i < n$, and let $V_i(\hat{x}_i) = \max\{\hat{x}_i - y, Z_i(\hat{x}_i)\}$. To derive the optimal stopping rule, we utilize an induction argument. We begin with the first step. If the seller has met with all n buyers, the decision is trivial: the good is allocated to the agent (either one of the buyers or the seller) with valuation \hat{x}_n . The expected surplus is $V_n(\hat{x}_n) = \hat{x}_n - y$.

Working backward, consider the planner's problem when the seller has met with only $n - 1$ of the n buyers. If the planner instructs the seller to stop meeting with buyers, again the good is allocated to the agent with valuation \hat{x}_{n-1} , yielding surplus $\hat{x}_{n-1} - y$. Alternatively, if the seller

meets with the next buyer, the expected surplus is $Z_{n-1}(\hat{x}_{n-1}) - y$, where

$$Z_{n-1}(\hat{x}) - y = -k + \int_{\underline{x}}^{\bar{x}} \max \{V_n(\hat{x}), V_n(x)\} f(x) dx,$$

so that

$$Z_{n-1}(\hat{x}) = -k + \hat{x}F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} xf(x) dx. \quad (25)$$

Notice immediately that $Z_{n-1}(x^*) = x^*$ and $Z'_{n-1}(\hat{x}) = F(\hat{x}) \in (0, 1)$, so that clearly x^* is the optimal cutoff after meeting with $n - 1$ buyers, and

$$V_{n-1}(\hat{x}) = \begin{cases} \hat{x} - y & \text{for } \hat{x} \geq x^* \\ Z_{n-1}(\hat{x}) - y & \text{for } \hat{x} < x^*. \end{cases}$$

Now consider the planner's problem after the seller has met with $n - 2$ buyers. We establish three important properties of $Z_{n-2}(\hat{x})$: (1) $Z_{n-2}(\hat{x}) = Z_{n-1}(\hat{x})$ for all $\hat{x} \geq x^*$; (2) $Z_{n-2}(\hat{x}) > \hat{x}$ for all $\hat{x} < x^*$; and (3) $\lim_{\hat{x} \rightarrow x^*} Z_{n-2}(\hat{x}) = Z_{n-1}(x^*) = x^*$. Given these three properties, along with the fact that $Z'_{n-2}(\hat{x}) = Z'_{n-1}(\hat{x}) = F(\hat{x}) \in (0, 1)$ for $x \geq x^*$, it follows immediately that $\hat{x} \geq Z_{n-2}(\hat{x})$ if, and only if, $\hat{x} \geq x^*$, and hence x^* is the optimal cutoff again.

After meeting with $n - 2$ buyers, the expected surplus from another meeting when $\hat{x}_{n-2} = \hat{x}$ is

$$Z_{n-2}(\hat{x}) - y = -k + V_{n-1}(\hat{x})F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} V_{n-1}(x)f(x) dx. \quad (26)$$

If $\hat{x} \geq x^*$, then $V_{n-1}(\hat{x}) = \hat{x} - y$ and thus $Z_{n-2}(\hat{x}) = Z_{n-1}(\hat{x})$. Alternatively, if $\hat{x} < x^*$, then $V_{n-1}(\hat{x}) = Z_{n-1}(\hat{x}) - y > \hat{x} - y$ and

$$\begin{aligned} Z_{n-2}(\hat{x}) &= -k + Z_{n-1}(\hat{x})F(\hat{x}) + \int_{\hat{x}}^{x^*} Z_{n-1}(x)f(x) dx + \int_{x^*}^{\bar{x}} xf(x) dx \\ &> -k + \hat{x}F(\hat{x}) + \int_{\hat{x}}^{x^*} xf(x) dx + \int_{x^*}^{\bar{x}} xf(x) dx = Z_{n-1}(\hat{x}) > \hat{x}. \end{aligned}$$

Finally, note that

$$\begin{aligned} \lim_{\hat{x} \rightarrow x^*} Z_{n-2}(\hat{x}) &= -k + Z_{n-1}(x^*)F(x^*) + \int_{x^*}^{\bar{x}} xf(x) dx \\ &= -k + x^*F(x^*) + \int_{x^*}^{\bar{x}} xf(x) dx = Z_{n-1}(x^*) = x^*. \end{aligned}$$

Therefore, the optimal cutoff after meeting with $n - 2$ buyers is x^* .

To summarize, we have established that the following is true for $j' = 2$: (i) $Z_{n-j'}(\hat{x}) =$

$Z_{n-j'+1}(\hat{x})$ for all $\hat{x} \geq x^*$; (ii) $Z_{n-j'}(\hat{x}) > \hat{x}$ for all $\hat{x} < x^*$; and (iii) $\lim_{\hat{x} \rightarrow x^*} Z_{n-j'}(\hat{x}) = Z_{n-j'+1}(x^*) = x^*$. It follows that

$$V_{n-j'}(\hat{x}) = \begin{cases} \hat{x} - y & \text{for } \hat{x} \geq x^* \\ Z_{n-j'}(\hat{x}) - y & \text{for } \hat{x} < x^*. \end{cases}$$

Now, suppose this is true for all $j' \in \{2, 3, \dots, j\}$. We will establish that it is also true for $j + 1$. After meeting with $n - j - 1$ buyers, the expected surplus from another meeting when $\hat{x}_{n-j-1} = \hat{x}$ is

$$Z_{n-j-1}(\hat{x}) - y = -k + V_{n-j}(\hat{x}) F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} V_{n-j}(x) f(x) dx. \quad (27)$$

If $\hat{x} \geq x^*$, then $V_{n-j}(\hat{x}) = \hat{x} - y$ and thus $Z_{n-j-1}(\hat{x}) = Z_{n-1}(\hat{x})$. Moreover, given the first assumption in the induction step, $Z_{n-j}(\hat{x}) = Z_{n-1}(\hat{x})$, so that $Z_{n-j-1}(\hat{x}) = Z_{n-j}(\hat{x})$.

Alternatively, if $\hat{x} < x^*$, then

$$\begin{aligned} Z_{n-j-1}(\hat{x}) &= -k + Z_{n-j}(\hat{x}) F(\hat{x}) + \int_{\hat{x}}^{x^*} Z_{n-j}(x) f(x) dx + \int_{x^*}^{\bar{x}} x f(x) dx \\ &> -k + \hat{x} F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} x f(x) dx = Z_{n-1}(\hat{x}) > \hat{x}. \end{aligned}$$

Finally, note that

$$\begin{aligned} \lim_{\hat{x} \rightarrow x^*} Z_{n-j-1}(\hat{x}) &= -k + Z_{n-j}(x^*) F(x^*) + \int_{x^*}^{\bar{x}} x f(x) dx \\ &= Z_{n-1}(x^*) = x^*. \end{aligned}$$

Therefore, we have that the optimal cutoff after meeting with $n - j - 1$ buyers is x^* .

Proof of Lemma 2

Substituting (2) into (6) and integrating by parts yields

$$S(x^*, \lambda) = x^* - y - \int_y^{x^*} q_0(x) dx.$$

Since x^* is independent of λ , the second derivative of $S(x^*, \lambda)$ with respect to λ then equals

$$\frac{d^2 S}{d\lambda^2} = - \int_y^{x^*} (1 - F(x))^2 q_0(x) dx < 0.$$

Proof of Proposition 1

Given Lemma 1, clearly a stopping rule of x^* is optimal at all sellers. Then, given Lemma 2, it follows that total surplus is maximized by assigning the same queue length $\lambda = \Lambda$ to all sellers.

Proof of Lemma 3

When the buyer declines to pay the asking price, he realizes that the auction will take place if and only if all other buyers visiting the same seller have a valuation below \tilde{x}^a . In the auction, the buyer will face a number of competitors which follows a Poisson distribution with mean $\lambda F(\tilde{x}^a)$. Their valuations will be distributed according to $F(x)/F(\tilde{x}^a)$.

Consider the case in which the buyer's valuation x is weakly below \tilde{x}^a . Standard arguments then imply that the expected payoff for the buyer from participating in the auction equals the integral of his trading probability.³⁴ That is,

$$\int_y^x e^{-\lambda F(\tilde{x}^a) \left(1 - \frac{F(x')}{F(\tilde{x}^a)}\right)} dx' = \int_y^x \frac{q_0(x')}{q_0(\tilde{x}^a)} dx', \quad (28)$$

where the integrand represents the probability that no other buyer has a valuation above x' , conditional on all valuations being below \tilde{x}^a . Multiplying this payoff by (7) yields the desired expression for $u(x; \tilde{x}^a)$. To derive the buyer's optimal bidding function $\hat{b}(x)$ in the first-price auction, note that (28) should be equal to the product of the probability that the buyer wins the auction and his payoff conditional on winning. That is,

$$\frac{q_0(x)}{q_0(\tilde{x}^a)} (x - \hat{b}(x)) = \int_y^x \frac{q_0(x')}{q_0(\tilde{x}^a)} dx',$$

where the first factor on the left-hand side represents the probability that no other buyer has a valuation above x , conditional on all valuations being below \tilde{x}^a . Solving for $\hat{b}(x)$ gives the desired result.

Next, consider the case in which the buyer's valuation x strictly exceeds \tilde{x}^a . When the buyer declines to pay the asking price, he maximizes his payoff in the auction by submitting a bid $\hat{b}(\tilde{x}^a)$. The expected payoff from this strategy is $x - \hat{b}(\tilde{x}^a)$ multiplied by the probability (7) that the auction will take place. Substitution of $\hat{b}(\tilde{x}^a)$ yields the desired expression for $u(x; \tilde{x}^a)$.

³⁴See Peters (2013) for a detailed discussion. Peters and Severinov (1997) explicitly analyze the case with a Poisson number of buyers.

Proof of a Unique Symmetric x^a

Consider a buyer with valuation x at a seller with asking price a and queue λ . Accepting the asking price gives the buyer a payoff $x - a$, while rejecting it yields $u(x; \tilde{x}^a)$. From equation (10), it readily follows that $0 < \frac{\partial u(x; \tilde{x}^a)}{\partial x} < 1$ since

$$\frac{\partial u(x; \tilde{x}^a)}{\partial x} = \begin{cases} \frac{\lambda(1-F(\tilde{x}^a))}{1-q_0(\tilde{x}^a)} q_0(x) & \text{if } x \leq \tilde{x}^a, \\ \frac{\lambda(1-F(\tilde{x}^a))}{1-q_0(\tilde{x}^a)} q_0(\tilde{x}^a) & \text{if } x > \tilde{x}^a. \end{cases} \quad (29)$$

For given a , λ and \tilde{x}^a , the buyer should therefore accept the asking price above a certain cutoff, which we denote—with a slight abuse of notation—by $x^a \equiv x^a(a, \lambda, \tilde{x}^a)$. If there is an interior cutoff, it is uniquely determined by

$$x^a - a = u(x^a; \tilde{x}^a); \quad (30)$$

otherwise, if a is too small or too large, respectively, the cutoff is $x^a = y$ or $x^a = \bar{x}$. Focusing on interior solutions, one can show that the cutoff is increasing in a , as the implicit function theorem and equation (29) imply

$$\frac{\partial x^a}{\partial a} = \left[1 - \frac{\partial u(x^a; \tilde{x}^a)}{\partial x^a} \right]^{-1} > 0. \quad (31)$$

Given symmetric strategies, $x^a = \tilde{x}^a$ needs to hold in equilibrium. To prove that a unique solution to this fixed point problem exists, consider how the (interior) cutoff x^a varies with \tilde{x}^a . Applying the implicit function theorem once more implies

$$\frac{\partial x^a}{\partial \tilde{x}^a} = \frac{\partial u(x^a; \tilde{x}^a)}{\partial \tilde{x}^a} \left[1 - \frac{\partial u(x^a; \tilde{x}^a)}{\partial x^a} \right]^{-1}.$$

By equation (29), the sign of this expression is determined by $\partial u(x^a; \tilde{x}^a) / \partial \tilde{x}^a$. Evaluated at $x^a = \tilde{x}^a$,

$$\left. \frac{\partial u(x^a; \tilde{x}^a)}{\partial \tilde{x}^a} \right|_{x^a = \tilde{x}^a} = -\lambda f(\tilde{x}^a) \frac{1 - q_0(\tilde{x}^a) - q_1(\tilde{x}^a)}{(1 - q_0(\tilde{x}^a))^2} \int_y^{\tilde{x}^a} q_0(x') dx' < 0.$$

Hence, the system of equations $x^a = \tilde{x}^a$ and (30) has a unique solution for every a and λ . By equation (31), this solution is increasing in a .

Proof of Proposition 2

Since the relation between a and x^a is one-to-one, given λ , the seller's maximization problem can be rewritten as a choice over x^a and λ , which turns out to be more convenient analytically.³⁵ Define $\widehat{R}(x^a, \lambda)$ as the revenue of a seller with asking price a , queue λ and cutoff $x^a \equiv x^a(a, \lambda)$. Substituting a and $\widehat{b}(x)$ into $R(a, \lambda)$, as given in (15), yields

$$\widehat{R}(x^a, \lambda) = x^a - (1 + \lambda) \int_y^{x^a} q_0(\widehat{x}) d\widehat{x} + \lambda \int_y^{x^a} F(x) q_0(x) dx.$$

One can derive $\widehat{U}(x^a, \lambda)$, i.e., the expected payoff of a buyer visiting this seller, in a similar fashion:

$$\widehat{U}(x^a, \lambda) = \frac{1}{\lambda} \left(\frac{1 - q_0(x^a)}{(1 - F(x^a))} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] + \lambda \int_y^{x^a} (1 - F(x)) q_0(x) dx \right). \quad (32)$$

The partial derivatives of $\widehat{R}(x^a, \lambda)$ are equal to

$$\begin{aligned} \frac{\partial \widehat{R}}{\partial x^a} &= 1 - Q_1(x^a) > 0 \\ \frac{\partial \widehat{R}}{\partial \lambda} &= \lambda \int_y^{x^a} (1 - F(x))^2 q_0(x) dx > 0 \end{aligned}$$

while the partial derivatives of $\widehat{U}(x^a, \lambda)$ are

$$\begin{aligned} \frac{\partial \widehat{U}}{\partial x^a} &= -\frac{1 - Q_1(x^a)}{\lambda} \left(1 - \frac{f(x^a)}{(1 - F(x^a))^2} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] \right) < 0 \\ \frac{\partial \widehat{U}}{\partial \lambda} &= -\frac{1 - Q_1(x^a)}{\lambda^2 (1 - F(x^a))} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] - \int_y^{x^a} q_0(x) (1 - F(x))^2 dx, \end{aligned}$$

since $Q_1(x^a) \equiv q_0(x^a) + q_1(x^a) < 1$. Therefore, the first-order conditions of the Lagrangian with

³⁵As noted in the text, a potential complication occurs if, for some deviation, there are multiple values of (x^a, λ) such that $\widehat{U}(x^a, \lambda) = \bar{U}$. The current formulation implicitly assumes that the seller calculates his payoffs from deviating using the most profitable pair (x^a, λ) . Naturally, if the deviation is not profitable for this pair, it is not profitable for any pair (x^a, λ) such that $\widehat{U}(x^a, \lambda) = \bar{U}$.

respect to x_a , λ , and μ , respectively, equal

$$0 = (1 - Q_1(x^a)) \left(1 - \frac{\mu}{\lambda} \left(1 - \frac{f(x^a)}{(1 - F(x^a))^2} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] \right) \right) \quad (33)$$

$$0 = \lambda \int_y^{x^a} q_0(x) (1 - F(x))^2 dx \left(1 - \frac{\mu}{\lambda} \right) - \frac{\mu}{\lambda^2} \frac{1 - Q_1(x^a)}{(1 - F(x^a))} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] \quad (34)$$

$$0 = \frac{1 - q_0(x^a)}{\lambda(1 - F(x^a))} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] + \int_y^{x^a} (1 - F(\hat{x})) q_0(\hat{x}) d\hat{x} - \bar{U}.$$

Solving (33) implies

$$\frac{\mu}{\lambda} = \left(1 - \frac{f(x^a)}{(1 - F(x^a))^2} \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] \right)^{-1},$$

so that (34) can be written as

$$0 = \left[\int_{x^a}^{\bar{x}} (x - x^a) dF(x) - k \right] \times \frac{\mu}{\lambda} \left\{ - \frac{\lambda f(x^a) \int_y^{x^a} q_0(x) (1 - F(x))^2 dx}{(1 - F(x^a))^2} - \frac{1 - Q_1(x^a)}{\lambda(1 - F(x^a))} \right\}.$$

Since the term in brackets on the second line of this equation is strictly negative, it must be that the unique solution for x^a satisfies $\int_{x^a}^{\bar{x}} (x - x^a) dF(x) = k$. From this, it immediately follows that $x^a = x^*$ and $\mu = \lambda = \Lambda$. Hence, the equilibrium is unique and it coincides with the solution to the planner's problem. Given $x^a = x^*$ and $\lambda = \Lambda$, the optimal asking price follows from (11).

Proof of Proposition 3

Consider a candidate equilibrium with market utility $0 < \bar{U} < \int_y^{\bar{x}} (x - y) f(x) dx - k$. Now, take an arbitrary mechanism m_1 that one or more sellers post in this equilibrium, which attracts a queue $\lambda_1 > 0$ that satisfies $\mathbf{U}(m_1, \lambda_1) = \bar{U}$. This mechanism yields the seller a payoff $\mathbf{R}(m_1, \lambda_1)$ and generates a surplus $\mathbf{S}(m_1, \lambda_1) = \mathbf{R}(m_1, \lambda_1) + \lambda_1 \mathbf{U}(m_1, \lambda_1) - y$.

Now consider an asking price mechanism m_a^* that implements the cutoff x^* and attracts a queue length $\lambda_a > 0$ that satisfies $\bar{U} = \mathbf{U}(m_a^*, \lambda_a)$.³⁶ From our results in Proposition 2, and the discussion

³⁶Note that λ_a depends on \bar{U} . However, in what follows, we suppress this implicit relationship for notational convenience.

that follows, we know that $\mathbf{S}(m_a^*, \lambda_a) = \mathbf{S}^*(\lambda_a)$ and

$$\left. \frac{\partial \mathbf{S}^*(\lambda)}{\partial \lambda} \right|_{\lambda=\lambda_a} = \bar{U}. \quad (35)$$

Since m_1 was chosen when m_a^* was feasible, it must be that

$$\mathbf{S}^*(\lambda_a) - \lambda_a \bar{U} + y = \mathbf{R}(m_a^*, \lambda_a) \leq \mathbf{R}(m_1, \lambda_1) = \mathbf{S}(m_1, \lambda_1) - \lambda_1 \bar{U} + y.$$

Next, let $\lambda^* = \arg \max_{\lambda} \mathbf{S}^*(\lambda) - \lambda \bar{U}$. Since $\mathbf{S}^*(\lambda)$ is strictly concave, we are assured of a unique solution that will satisfy

$$\left. \frac{\partial \mathbf{S}^*(\lambda)}{\partial \lambda} \right|_{\lambda=\lambda^*} = \bar{U}.$$

From (35), it must be that $\lambda^* = \lambda_a$. Moreover, since $\mathbf{S}(m_1, \lambda_1) \leq \mathbf{S}^*(\lambda_1)$, we can write

$$\begin{aligned} \mathbf{R}(m_1, \lambda_1) = \mathbf{S}(m_1, \lambda_1) - \lambda_1 \bar{U} + y &\leq \mathbf{S}^*(\lambda_1) - \lambda_1 \bar{U} + y \\ &\leq \mathbf{S}^*(\lambda^*) - \lambda^* \bar{U} + y \\ &= \mathbf{S}^*(\lambda_a) - \lambda_a \bar{U} + y = \mathbf{R}(m_a^*, \lambda_a). \end{aligned}$$

Hence, for any equilibrium mechanism m_1 , it must be that $\mathbf{R}(m_1, \lambda_1) = \mathbf{R}(m_a^*, \lambda_a)$. Moreover, it must also be that $\lambda_1 = \lambda_a = \lambda^*$. Otherwise, for any $\lambda_1 \neq \lambda^*$, $\mathbf{R}(m_1, \lambda_1) \leq \mathbf{S}^*(\lambda_1) - \lambda_1 \bar{U} + y < \mathbf{S}^*(\lambda^*) - \lambda^* \bar{U} + y = \mathbf{R}(m_a^*, \lambda^*)$. This concludes the proof of Proposition 3.

Finally, since these results hold for all sellers, it must be that $\lambda^* = \Lambda$ in equilibrium. Therefore, for any m_1 offered in equilibrium, we must $\lambda_1 = \Lambda$, $\mathbf{U}(m_1, \lambda) = \mathbf{U}(m_a^*, \Lambda) = \bar{U}^*$ and $\mathbf{R}(m_1, \lambda_1) = \mathbf{R}(m_a^*, \Lambda) = \mathbf{S}^*(\Lambda) - \Lambda \bar{U} + y$; this is the result reported in Corollary 2.

Proof of Lemma 4

Consider a seller with n buyers. The sale price will be above asking price if $i \in \{2, \dots, n\}$ buyers draw a valuation between a and x^a while the remaining $n - i$ buyers draw a valuation below a . The probability of this event equals

$$\sum_{i=2}^n \frac{n!}{i!(n-i)!} (F(x^a) - F(a))^i F(a)^{n-i} = F(x^a)^n - n(F(x^a) - F(a))F(a)^{n-1} - F(a)^n.$$

Taking the expectation over n yields the desired expression.

Proof of Lemma 5

For the planner's problem, the proof proceeds by induction, much like the proof of lemma 1. Suppose that n buyers visit a seller, the first $n - 1$ buyers learn their valuation, and no trade has taken place because $\hat{x}_{n-1} \equiv \max \{y, x_1, \dots, x_{n-1}\} < x^*$. In this case, the planner has two options: either let buyer n incur the inspection cost k and base the ensuing trading decision on \hat{x}_n , or avoid the inspection cost by instructing the seller to trade with buyer n without knowing his valuation.³⁷

In the former case, expected surplus generated by the match is $Z_{n-1}(\hat{x}_{n-1}) - y$, where

$$Z_{n-1}(\hat{x}) = -k + \hat{x}F(\hat{x}) + \int_{\underline{x}}^{\hat{x}} xf(x) dx,$$

while the latter case yields an expected surplus equal to $\int_{\underline{x}}^{\bar{x}} xf(x) dx - y$. Clearly, inspection is preferred if and only if $Z_{n-1}(\hat{x}_{n-1}) - y > \int_{\underline{x}}^{\bar{x}} xf(x) dx - y$, or equivalently

$$k < \int_{\underline{x}}^{\hat{x}_{n-1}} (\hat{x}_{n-1} - x) f(x) dx.$$

This condition needs to hold for any feasible value of \hat{x}_{n-1} in order to guarantee inspection by the last buyer. Since the right-hand side is strictly increasing in \hat{x}_{n-1} , (22) is a necessary and sufficient condition. The final step is then to show that this condition implies that inspection is also optimal after meeting with $n - j - 1$ buyers for $j \in \{1, \dots, n - 1\}$. This follows immediately from $Z_{n-j-1}(\hat{x}) \geq Z_{n-1}(\hat{x})$ for all \hat{x} and j , as shown in the proof of lemma 1.

Next, we analyze the market equilibrium described in section 4. Consider a deviating buyer who does not inspect the good and therefore does not know his valuation. This deviant has three options: 1) submit a bid below y , which will be rejected; 2) submit a bid between y and a^* ; or 3) bid the asking price and trade immediately. The optimal choice is to behave like a buyer who has a valuation equal to the unconditional expectation of x , which we denote by $x^e = E_F[x] \equiv \int_{\underline{x}}^{\bar{x}} xf(x) dx$. That is, he should submit a bid below y if $x^e < y$ and should bid $\hat{b}(x^e)$ if $x^e \in [y, x^*)$. Note that the remaining case, $x^e \in [x^*, \bar{x}]$, cannot occur under (22), since it implies

$$\begin{aligned} x^e &= - \int_{\underline{x}}^{x^*} (x^* - x) f(x) dx + k + x^* \\ &< - \int_{\underline{x}}^y (y - x) f(x) dx + k + x^* < x^*. \end{aligned}$$

To see whether the deviant benefits from not inspecting the good, define an auxiliary distribu-

³⁷The planner can of course also instruct the seller to immediately trade with the agent with valuation \hat{x}_{n-1} , but, as shown in lemma 1, this is dominated by learning the valuation of the last buyer since $\hat{x}_{n-1} < x^*$.

tion $\tilde{F}(x)$ that resembles $F(x)$, except that the mass below y and above x^* is concentrated as mass points at y and x^* , respectively. That is,

$$\tilde{F}(x) = \begin{cases} 0 & \text{if } x < y \\ F(x) & \text{if } y \leq x \leq x^* \\ 1 & \text{if } x^* < x. \end{cases}$$

Let $\tilde{x}^e = E_{\tilde{F}}[x]$ denote the expectation of x under this modified distribution. Under (22), it then follows that $\tilde{x}^e > x^e$, since

$$\begin{aligned} \tilde{x}^e - x^e &= \int_{\underline{x}}^y (y - x) dF(x) - \int_{x^*}^{\bar{x}} (x - x^*) dF(x) \\ &= \int_{\underline{x}}^y (y - x) dF(x) - k > 0. \end{aligned}$$

The fact that $u(x; x^*)$ is an increasing function of x then implies that $u(x^e; x^*) < u(\tilde{x}^e; x^*)$, while the convexity of $u(x; x^*)$ in x implies that $u(\tilde{x}^e; x^*) < E_{\tilde{F}}[u(x; x^*)]$ by Jensen's inequality. Note, however, that $E_{\tilde{F}}[u(x; x^*)]$ exactly equals the payoff from inspection, since

$$\begin{aligned} E_{\tilde{F}}[u(x; x^*)] &= u(y; x^*) \tilde{F}(y) + \int_y^{x^*} u(x; x^*) d\tilde{F}(x) + (1 - \tilde{F}(x^*)) u(x^*; x^*) \\ &= \int_y^{x^*} u(x; x^*) dF(x) + (1 - F(x^*)) u(x^*; x^*) \\ &= \frac{\lambda(1 - F(x^*))}{1 - q_0(x^*)} \left[\int_y^{x^*} \int_y^x q_0(\tilde{x}) d\tilde{x} dF(x) + (1 - F(x^*)) \int_y^{x^*} q_0(\tilde{x}) d\tilde{x} \right] \end{aligned}$$

Hence, the payoff from inspection is strictly higher than the payoff from not inspecting.

To show that equation (22) is also necessary, consider the limit $\Lambda \rightarrow 0$, such that a buyer who meets with a seller knows that, with probability 1, he does not face competition from other buyers. The optimal bid in that case is y and it follows immediately that inspection is better only if equation (22) holds.

References

- Acemoglu, D. and Shimer, R. (1999). Efficient unemployment insurance. *Journal of Political Economy*, 107:893–928.
- Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2012). A note on Peters and Severinov, “Competition among sellers who offer auctions instead of prices”. *Journal of Economic Theory*, 147:389–392.
- Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2016). Directed search in the housing market. *Review of Economic Dynamics*, 19:218–231.
- Arnold, M. A. (1999). Search, bargaining and optimal asking prices. *Real Estate Economics*, 27(3):453–481.
- Budish, E. B. and Takeyama, L. N. (2001). Buy prices in online auctions: Irrationality on the internet? *Economics Letters*, 72:325–333.
- Bulow, J. and Klemperer, P. (2009). Why do sellers (usually) prefer auctions? *American Economic Review*, 99(4):1544–1575.
- Burdett, K., Shi, S., and Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109:1060–1085.
- Burguet, R. (1996). Optimal repeated purchases when sellers are learning about costs. *Journal of Economic Theory*, 68(2):440–455.
- Burguet, R. and Sákovic, J. (1999). Imperfect competition in auction designs. *International Economic Review*, 40(1):231–247.
- Chen, Y. and Rosenthal, R. W. (1996). Asking prices as commitment devices. *International Economic Review*, 37(1):129–155.
- Crémer, J., Spiegel, Y., and Zheng, C. Z. (2009). Auctions with costly information acquisition. *Economic Theory*, 38(1):41–72.
- Delacroix, A. and Shi, S. (2013). Pricing and signaling with frictions. *Journal of Economic Theory*, 148:1301–1332.
- Eeckhout, J. and Kircher, P. (2010). Sorting vs screening - prices as optimal competitive sales mechanisms. *Journal of Economic Theory*, 145:1354–1385.
- Engelbrecht-Wiggans, R. (1987). On optimal reservation prices in auctions. *Management Science*, 33(6):763–770.
- Hosios, A. J. (1990). On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies*, 57(2):279–298.
- Julien, B., Kennes, J., and King, I. (2000). Bidding for labor. *Review of Economic Dynamics*, 3(4):619–649.
- Kim, K. and Kircher, P. (2015). Efficient competition through cheap talk: The case of competing auctions. *Econometrica*, 83(5):1849–1875.

- Lester, B. (2015). Why ask? the role of asking prices in transactions. *Federal Reserve Bank of Philadelphia Business Review*, 4:1–4.
- Lester, B., Visschers, L., and Wolthoff, R. (2013). Competing with asking prices. *Federal Reserve Bank of Philadelphia Working Paper Series*, (13-7).
- Lester, B., Visschers, L., and Wolthoff, R. (2015). Meeting technologies and optimal trading mechanisms in competitive search markets. *Journal of Economic Theory*, 155:1–15.
- Levin, D. and Smith, J. L. (1994). Equilibrium in auctions with entry. *The American Economic Review*, pages 585–599.
- Lippman, S. A. and McCall, J. J. (1976). The economics of job search: A survey. *Economic Inquiry*, XIV:155–189.
- Mathews, T. (2004). The impact of discounting on an auction with a buyout option: a theoretical analysis motivated by eBay’s buy-it-now feature. *Journal of Economics*, 81(1):25–52.
- McAdams, D. (2015). On the benefits of dynamic bidding when participation is costly. *Journal of Economic Theory*, 157:959–972.
- McAfee, R. P. (1993). Mechanism design by competing sellers. *Econometrica*, 61(6):pp. 1281–1312.
- McAfee, R. P. and McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, 25:699–738.
- McAfee, R. P. and McMillan, J. (1988). Search mechanisms. *Journal of Economic Theory*, 44:99–123.
- Menzio, G. (2007). A theory of partially directed search. *Journal of Political Economy*, 115:748–769.
- Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105:385–411.
- Morgan, P. and Manning, R. (1985). Optimal search. *Econometrica*, 53:923–944.
- Mortensen, D. T. (1982). Efficiency of mating, racing, and related games. *American Economic Review*, 72:968–979.
- Pancs, R. (2013). Sequential negotiations with costly information acquisition. *Games and Economic Behavior*, 82:522–543.
- Peters, M. (2001). Surplus extraction and competition. *The Review of Economic Studies*, 68(3):613–631.
- Peters, M. (2013). *The Handbook of Market Design*, chapter Competing Mechanisms. Oxford University Press.
- Peters, M. and Severinov, S. (1997). Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75:141–179.
- Reynolds, S. S. and Wooders, J. (2009). Auctions with a buy price. *Economic Theory*, 38(1):9–39.

- Schmitz, P. W. (2003). On second-price auctions and imperfect competition. *Journal of Mathematical Economics*, 39:901–909.
- Stacey, D. (2012). Information, commitment, and separation in illiquid housing markets. mimeo.
- Virág, G. (2010). Competing auctions: finite markets and convergence. *Theoretical Economics*, 5(2):241–274.
- Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–654.

Figure 1: Optimal Mechanisms




Figure 2: Transaction Price CDF






Figure 3: Effect of an Increase in the Buyer/Seller Ratio




solid blue line = original equilibrium; dashed orange line = equilibrium with higher Λ .

Figure 4: Effect of a Decrease in the Inspection Cost
CDF



solid blue line = original equilibrium; dashed orange line = equilibrium with lower k .

Figure 5: Effect of Technological Improvement
CDF



solid blue line = original equilibrium; dashed orange line = equilibrium with more information.

Figure 6: Transaction Prices Above the Asking Price
CDF

