

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

“Competing with Asking Prices”

Benjamin LESTER

Federal Reserve Bank of Philadelphia

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. Despite their prevalence in a variety of real world markets, asking prices have received little attention in the academic literature. We construct an environment with a few simple, realistic ingredients and demonstrate that using an asking price is optimal: it is the pricing mechanism that maximizes sellers' revenues and it implements the efficient outcome in equilibrium. We provide a complete characterization of this equilibrium and use it to explore the positive implications of this pricing mechanism for transaction prices and allocations.

Friday, March 6, 2015, 10:30-12:00
Room 126, Extranef building at the University of Lausanne

Competing with Asking Prices*

Benjamin Lester[†]

Federal Reserve Bank of Philadelphia

Ludo Visschers[‡]

University of Edinburgh
& Universidad Carlos III, Madrid

Ronald Wolthoff

University of Toronto

October 24, 2013

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. Despite their prevalence in a variety of real world markets, asking prices have received little attention in the academic literature. We construct an environment with a few simple, realistic ingredients and demonstrate that using an asking price is optimal: it is the pricing mechanism that maximizes sellers' revenues and it implements the efficient outcome in equilibrium. We provide a complete characterization of this equilibrium and use it to explore the positive implications of this pricing mechanism for transaction prices and allocations.

Keywords: Asking Prices, Directed Search, Inspection Costs, Efficiency

JEL codes: C78, D44, D82, D83, L11, R31, R32

*We would like to thank Jim Albrecht, Jan Eeckhout, Andrea Galeotti, Pieter Gautier, Lu Han, Benoit Julien, Philipp Kircher, Guido Menzio, Peter Norman, József Sákovics, Shouyong Shi, Andy Skrzypacz, Xianwen Shi, Rob Shimer, Susan Vroman, Gábor Virág, and Randy Wright for helpful comments. All errors are our own.

[†]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. This paper is available free of charge at www.philadelphiafed.org/research-and-data/publications/working-papers/.

[‡]Ludo Visschers gratefully acknowledges financial support from the Juan de la Cierva Grant; project grant ECO2010-20614 (Dirección general de investigación científica y técnica), and the Bank of Spain's Programa de Investigación de Excelencia.

“We are eager to hear ... about businesses that meet all of the following criteria:

(1) Large purchases; (2) Demonstrated consistent earning power; (3) Businesses earning good returns on equity while employing little or no debt; (4) Management in place; (5) Simple businesses; **(6) An offering price (we don’t want to waste our time or that of the seller by talking, even preliminarily, about a transaction when price is unknown)... We don’t participate in auctions.**”¹

– Berkshire Hathaway Inc., “Acquisition Criteria,” 2012 Annual Report

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 such an offer could potentially be accepted if the seller receives no better offers.² Though asking prices are prevalent in a variety of real world 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; that is, 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 use it to explore the positive 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, anybody who has recently purchased a house or a car, rented an apartment, walked through a bazaar, or perused the classifieds knows that an asking price seems to play a prominent role in the sale of many goods (and services). The question is: how and why can this mechanism be optimal?

¹Bold added for emphasis; italics present in original.

²What we call an asking price goes by several other names as well, including an “offering price” (as in the epigraph), a “list price” (used in the sale of houses and cars), or a “buy-it-now” or “take-it” price (used 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.

Loosely speaking, our answer requires two ingredients. First, the sellers in our environment *compete for buyers*. That is, 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 responsive to the mechanism they post. 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.³

In an environment with these two ingredients, a natural tension arises. *Ceteris paribus*, sellers would like to place no limit on the number of buyers who inspect their good or on the offers that these buyers make. However, such a pricing mechanism is not particularly attractive to buyers. Instead, when a buyer incurs the inspection cost, he wants to be assured that he has a reasonable chance of actually acquiring the good (i.e., that another buyer hasn't already inspected the good and discovered a very high valuation). We show that the asking price mechanism provides this assurance by implementing a stopping rule, so that the good is allocated to the first buyer who has a sufficiently high valuation, and spares the remaining potential buyers from inspecting the good "in vain." In other words, echoing the sentiments of Berkshire Hathaway's chairman Warren Buffett in the epigraph above, the asking price mechanism in our model *serves as a promise by sellers not to waste the buyers' time and energy*. In a competitive setting, this promise is the most effective way for a seller to attract buyers and thus, in equilibrium, sellers use this pricing mechanism. 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,

³For example, suppose the goods are houses that are roughly equivalent along easily describable dimensions (size, general location, and so on). However, each home has idiosyncratic features that may make them more or less attractive to every prospective buyer, and these features are revealed upon inspection; e.g., an individual who likes to cook will want to examine a home's kitchen. Quite often, learning one's true valuation may require more research than is afforded by a quick tour; e.g., an individual who needs to build a home office may want to bring in an architect to get an estimate of how much it will cost. All of these activities are costly, either because they take time or because they require explicit costs (like hiring an architect). Similar costs exist for purchasing a car, renting an apartment, or even hiring an accountant. Perhaps surprisingly, these costs can even be significant for buyers purchasing goods on websites such as eBay or Craigslist, as documented by Bajari and Hortacsu (2003).

and buyers can observe what each seller posts and visit the seller that offers the highest expected payoff. The matching 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 matching 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. This strategy of “sequential search with recall” is optimal because it balances the losses associated with additional buyers incurring the inspection cost against the gains associated with finding a buyer who values the good more than all of the previous buyers. 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, characterize the trading mechanism that emerges in equilibrium, and study the efficiency properties of this equilibrium. 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.⁴ 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

⁴As we discuss below, we consider the information made available to buyers to be a feature of the pricing mechanism.

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 such a simple, realistic instrument as 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* price — without any other devices such as meeting fees, bidding subsidies, or reserve prices — that controls both margins; it determines 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 there exists an alternative mechanism that could potentially increase sellers' profits. To address this question, in Section 5 we consider a similar game to the one described above, but we allow sellers to select a mechanism from an arbitrary set of pricing schemes. In this more general environment, we establish that an asking price mechanism maximizes a seller's payoff, irrespective 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, there is no equilibrium in which sellers earn higher expected payoffs 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 are a robust and compelling way to deal with the frictions in our environment.

The normative analysis described above leaves us with a tractable, micro-founded theory of asking prices, which we think could be a useful benchmark for both theoretical and empirical work that focuses on markets in which this pricing mechanism is prevalent. In Section 6, we flesh out just a few of the model's positive implications for a variety of observable outcomes. In particular, we study the level of asking prices set by sellers and the corresponding distribution of transaction prices that occur in equilibrium. We examine how these variables change with 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. Moreover, we analyze the relationship between the expected transaction price and the number of buyers who inspect the good before trade; to the extent that each inspection takes time, this analysis provides new insights into the relationship

between transaction prices and time on the market.

In Section 7, we discuss a few of our key assumptions, along with several interesting extensions of our basic framework. A number of these extensions could be particularly helpful for taking our model to data. For one, we show 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 confronting, e.g., house prices.⁵ We also discuss how allowing for heterogeneity across buyers could lead to endogenous market segmentation, where each sub-market has a different asking price, a different market tightness, and subsequently produces different patterns for prices and allocations. Finally, 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 an uncomplicated, realistic instrument — the asking price — is both revenue-maximizing and efficient in an environment with two simple ingredients. The second is positive: we provide a rationale for the use of asking prices — namely, as a promise by sellers not to waste buyers’ time and resources — and explore the empirical implications of this mechanism 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 a random number of 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 have often found 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 is independent of the pricing mechanism. Hence, a simple mechanism, like a price or a wage, 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),

⁵Case and Shiller (2003) and Han and Strange (2013) document housing prices and show that a significant fraction of transactions occur at the asking price, and (in some locations) a number of transactions also occur above the asking price.

Peters and Severinov (1997), Burguet and Sákovics (1999), Eeckhout and Kircher (2010), Virág (2010), Kim and Kircher (2012), and Albrecht et al. (2012b), 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. We establish that the mechanism that achieves the efficient ex post allocation also induces efficient search decisions by buyers ex ante. Moreover, this sequential mechanism is surprisingly simple: it does *not* require that sellers have access to rarely observed devices like meeting fees or bidding subsidies, which allow sellers to transfer or appropriate part of the expected surplus without affecting its size.

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. However, the prevailing wisdom in this literature is that the optimal mechanism in this environment will require that the seller have access to a number of different instruments. For example, Crémer et al. (2009) show that, in this type of environment, 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.⁶ Deeming such a mechanism implausibly elaborate and demanding, Bulow and Klemperer (2009) instead focus on two “plain vanilla” mechanisms — a simultaneous auction and a simple sequential mechanism — and identify a trade-off between revenue-maximization and efficiency. In contrast to these papers, the mechanism we consider is remarkably simple (sellers post a single number) and yet the trade-off in Bulow and Klemperer (2009) does not arise; the asking price is both revenue-maximizing and efficient.⁷

The reason our results differ from those in the existing literature is twofold. The first crucial difference 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 endogenous. Hence, the sellers in our model have an incentive to post mechanisms that implement the efficient allocation, since they are the

⁶These take-it-or-leave-it offers act like a sequence of asking prices which, along with participation fees, are calibrated just right to extract all of the buyers’ surplus. In a similar environment, Burguet (1996) also shows that one way to implement the efficient allocation is through a relatively complicated sequence of reserve prices and bidding subsidies.

⁷Roberts and Sweeting (2012) also compare the performance of a simultaneous auction and a simple sequential mechanism, though they assume buyers have some information about their valuation before making an entry decision. Interestingly, 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.

residual claimant on any extra surplus that they create.⁸ The second crucial difference between our environment and that of, e.g., Crémer et al. (2009) and Bulow and Klemperer (2009) is that these papers assume that prospective buyers observe both the number and behavior of previous bidders who have entered the mechanism, whereas we allow the sellers in our model to choose what information is available. Importantly, they choose not to reveal information about previous bidders. Instead, when the buyers in our model inspect the good, they only know that no previous bidder has discovered a valuation high enough to bid the asking price. As we discuss in greater detail in Section 4, this (endogenously determined) information structure is one important reason that the trade-off in Bulow and Klemperer (2009) disappears in our environment.

Turning to the positive results, our explanation is of course not the only plausible reason why asking prices might serve a useful role. For one, if buyers are risk averse, an asking price offers a way of reducing the uncertainty an individual buyer faces in certain (not necessarily optimal) types of auctions, and hence introducing this mechanism can potentially increase a seller's revenues; see, e.g., Budish and Takeyama (2001), Mathews (2004), or Reynolds and Wooders (2009). A second explanation for asking prices, which also assumes that it is costly for buyers to learn their valuation, is proposed by Chen and Rosenthal (1996) and Arnold (1999). In their environment, a holdup problem emerges when a buyer and seller bargain over the terms of trade after the buyer incurs the inspection cost. An asking price is treated as a ceiling on the bargaining outcome and thus partially solves this holdup problem. This explanation is close to ours in spirit, in the sense that an asking price serves as an *ex ante* guarantee that some rents will be transferred from the seller to the buyer. However, in contrast to our work, the asking price in these papers has no allocative role, nor is it clear that the asking price is the most efficient way of solving this holdup problem.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, the price that sellers advertise 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. Hence, this quoted 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

⁸This insight is not only true in the competitive search models cited above, but also in a very related literature that studies a single seller who selects the mechanism by which they want to sell their good, and a number of buyers who can choose to participate in the mechanism; see Engelbrecht-Wiggans (1987), McAfee and McMillan (1987), and Levin and Smith (1994) for seminal contributions. The difference is that the “cost” to a buyer from visiting a seller in the former literature is the opportunity cost of not visiting other sellers (which is an equilibrium object), while the cost of participating in the latter is an explicit entry cost (which is taken as exogenous).

⁹More generally, these theories of asking prices 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.

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. (2012a), 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, so that $\Lambda = \theta_b/\theta_s$ 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.

Matching. Buyers can visit a single seller in attempt to trade, but the matching process is frictional. In particular, buyers cannot coordinate with one another when choosing a seller to visit, and hence the number of buyers to arrive at each seller, n , will be stochastic. As is customary in the literature on directed (or competitive) search, we assume that n is distributed according to the Poisson distribution with parameter λ , which represents the expected number or "queue length" of buyers to arrive at a particular seller.¹² 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 continuous support on the interval $[\underline{x}, \bar{x}]$, and that the realization of x is the buyer's private information.

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 (2012) also study the signalling role of prices in directed search equilibria. However, in papers like Albrecht et al. (2012a) 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.

¹²The Poisson distribution is commonly used in these models because there are explicit micro-foundations: one can study a game with a finite number of agents in which each buyer chooses a single seller, but buyers are restricted to symmetric strategies. The matching technology that emerges—*urn-ball* matching—converges to a Poisson matching function as the number of buyers and sellers gets large. See, e.g., Burdett et al. (2001).

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. However, we stress that much of the analysis remains similar when $y < \underline{x}$, though the algebra is slightly more involved.

Inspection Costs. A key friction in the model is that the process of inspecting a good is costly to the buyer. In particular, 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. For instance, when a buyer is looking to purchase a car, it is costly for him to take time away from other productive activities to sit down with the seller, go over the car, take it for a test drive, and perhaps take it to a mechanic to be inspected. We use k to capture all of these costs, both implicit and explicit.

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.

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 matching 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. Working backward, we begin by characterizing these trading rules at an arbitrary seller, taking the queue length as given, and return later to characterize the optimal assignment of queue lengths across sellers.

Optimal Trading Protocol. Suppose n buyers arrive at a seller. The first result that we present in this section is that it is optimal for the planner to learn the buyers' valuations sequentially (i.e., one at a time), preserving the option to stop after each inspection $i \leq n$ and allocate the good to either one of the i buyers who has inspected the good, or to instruct the seller to retain the good for himself. Intuitively, this *sequential search strategy* is optimal because it allows the decision of whether an additional buyer should incur the inspection cost to be contingent on the realization of previous buyers' valuations; the planner can economize on k if the expected gain from learning the valuation of one or more additional buyers is small. The proof of this result is straightforward and follows very closely to that in Morgan and Manning (1985).

Lemma 1. *The optimal strategy for the planner is a sequential search strategy.*

We now characterize the planner's optimal stopping rule. In general, after a seller has met with i buyers with valuations (x_1, \dots, x_i) , when $n \geq i$ total buyers have arrived, the planner's decision rule is a function that dictates whether a seller should (a) keep meeting with additional buyers (if there are any), (b) trade with buyer $i' \in \{1, \dots, i\}$, or (c) retain the good for his own consumption. However, several features of the environment simplify this function immediately. First, notice that the only relevant information in the vector (x_1, \dots, x_i) is the maximum valuation; clearly, if the planner instructs the seller to stop meeting buyers, the good must be allocated to either the buyer with the highest valuation or the seller himself. Given this, it will be convenient to denote by $\hat{x}_i \equiv \max\{y, x_1, \dots, x_i\}$. Second, while the cost of an additional meeting between the seller and a buyer is a constant (k), we conjecture and later confirm that the continuation value of meeting with the $(i + 1)^{th}$ buyer is increasing in \hat{x}_i , so that the planner's optimal decision rule is a cutoff strategy. Let us denote the planner's cutoff by $x_{i,n}^p$ for $n \in \mathbb{N} \setminus \{1\}$ and $i \in \{1, 2, \dots, n - 1\}$, so that the seller will stop meeting buyers and the good will be consumed by the agent with valuation \hat{x}_i if, and only if, $\hat{x}_i \geq x_{i,n}^p$. Otherwise, if $\hat{x}_i < x_{i,n}^p$, the seller will continue to meet with the next buyer. Below we establish that, in fact, the optimal stopping rule is *stationary*; in particular, it does not depend on the number of buyers who have already inspected the good, i , or the number of buyers still available to inspect the good, $n - i$.¹³

¹³See Lippman and McCall (1976) for a similar result. It is worth noting that the stationarity of $x_{i,n}^p$ implies that whether or not the planner can observe i and/or n is ultimately immaterial, as he would choose the same cutoff rule in either environment.

Lemma 2. 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 the social surplus implementing the following 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 if, and only if, $\hat{x}_i \geq x_{i,n}^p = x^*$. 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 .

The proof of Lemma 2 utilizes an induction argument; we sketch the first step here, for intuition, and present the formal proof in the appendix. Suppose an arbitrary number $n \geq 2$ buyers arrive at a seller, and let $Z_{n-j,n}(\hat{x}) - y$ denote the net expected surplus from continuing to learn buyers' valuations, given that the maximum valuation of the seller and the $n - j$ buyers sampled so far is \hat{x} , for some $j \in \{1, 2, \dots, n - 1\}$. For notational convenience, we will drop the second subscript "n", so that $Z_{n-j,n}(\hat{x}) \equiv Z_{n-j}(\hat{x})$, $x_{n-j,n}^p \equiv x_{n-j}^p$, and so on; it should be understood that the analysis is for an arbitrary value of n .

Figure 1 plots $Z_{n-1}(\hat{x})$, $Z_{n-2}(\hat{x})$, and \hat{x} . Notice three facts: (i) $Z_{n-1}(\hat{x}) > \hat{x}$ if and only if $\hat{x} < x^*$; (ii) $Z_{n-2}(\hat{x}) > Z_{n-1}(\hat{x})$ if and only if $\hat{x} < x^*$; and (iii) $Z_{n-2}(\hat{x}) = Z_{n-1}(\hat{x})$ for $\hat{x} \geq x^*$. The first of these facts is by construction: x^* is defined as the value of \hat{x} that makes the planner indifferent between learning the valuation of one additional buyer and stopping. Hence, it should be obvious that $x_{n-1}^p = x^*$ is optimal.

INSERT FIGURE 1 HERE

Now consider the second and third facts. Intuitively, $Z_{n-2}(\hat{x}) - Z_{n-1}(\hat{x}) \geq 0$ captures the option value of being able to sample two more buyers instead of only one. However, if $\hat{x}_{n-2} \geq x^* = x_{n-1}^p$, then this option is never exercised after $n - 1$ valuations have been discovered, and thus $Z_{n-2}(\hat{x}) = Z_{n-1}(\hat{x})$; since the n^{th} buyer is never sampled, both $Z_{n-2}(\hat{x})$ and $Z_{n-1}(\hat{x})$ denote the expected surplus from sampling just one more buyer. Alternatively, if $\hat{x}_{n-2} < x^* = x_{n-1}^p$, then the option to sample two more buyers instead of one *can* be valuable, but this additional value is realized only in the event that $x_{n-1} \in (\hat{x}_{n-2}, x^*)$: if $x_{n-1} < \hat{x}_{n-2}$ then $\hat{x}_{n-2} = \hat{x}_{n-1}$, and if $x_{n-1} > x^*$ then the n^{th} buyer is never sampled. Hence, as $\hat{x}_{n-2} \rightarrow x^*$, this option value converges to zero and $Z_{n-2}(\hat{x}) \rightarrow Z_{n-1}(\hat{x})$. As a result, clearly $Z_{n-2}(\hat{x}) > \hat{x}$ if and only if $\hat{x} < x^*$; that is, $x_{n-2}^p = x^*$ is optimal. This completes the sketch of the first induction step; the remainder proceeds in much the same way.

Optimal Queue Lengths. We turn now to the optimal assignment of queue lengths across sellers, given the planner's decision rule for after buyers arrive. Notice immediately that the planner's optimal cutoff is not only independent of n and i , it is 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^p .

Let $S(x^p, \lambda)$ denote the expected surplus generated at an individual seller who is assigned a cutoff x^p and a queue length λ . In order 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 who 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 from below as x^p goes to \bar{x} . 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 $\int_0^{\theta_s} \lambda_j dj = \theta_b$. We show in the appendix that $S(x^*, \lambda)$ is strictly concave in its second argument.¹⁴ Since x^* is independent of λ , this is 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.*

The fact that the efficient cutoff is invariant to the number of buyers who have arrived at a seller, or to the number of buyers that have already inspected the good, is suggestive that an equally simple device may be able to implement the solution to the planner's problem in a decentralized economy. This is the focus of the next section.

¹⁴Two factors contribute to the concavity of S 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.

4 The Decentralized Equilibrium

We now consider the decentralized equilibrium. In this setting, an obvious challenge to implementing the social planner's allocation is that the seller does not observe – and hence cannot condition his decision on – a buyer's private valuation. However, we introduce a pricing mechanism, which we call an “asking price mechanism,” which implements a stopping rule for a seller who meets with a sequence of buyers, while still allowing for the recall of previous meetings. We characterize optimal asking prices and show that the decentralized equilibrium coincides with the solution to the planner's problem. Moreover, while the asking price mechanism we propose may at this point seem somewhat arbitrary, we establish below that, in fact, this is the sellers' *optimal* mechanism; that is, there exists no equilibrium mechanism that provides sellers with a higher payoff.

Asking Price Mechanism. The trading process with asking prices proceeds as follows. First, each seller posts an asking price, which we denote by a . Buyers observe all asking prices and decide which seller to visit, taking as given the decisions of other buyers. This determines the queue length λ at each seller. As in the planner's problem, the number of buyers to arrive at each seller is then a random variable n distributed according to the Poisson distribution with parameter λ . At each seller, after the realization of n , all buyers are placed in a random order, and the seller meets with the first buyer in the queue. The buyer incurs the inspection cost k and learns his valuation. After this, the buyer submits a bid b . Buyers know neither n nor their place in the queue.¹⁵

If $b \geq a$, the 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 $b < y$, then the bid is rejected. 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 seller shows the buyer his good, whereupon the 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 .

In sum, 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.

¹⁵Note 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.

Buyer's Bidding Function. Working backward, we begin by characterizing the optimal bidding strategy of 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 .

Let us denote by $b(x)$ the optimal bid of a buyer who meets with the seller and draws valuation x . To characterize $b(x)$, we conjecture and then confirm several important properties. First, it is straightforward to establish that a buyer should bid $b(x) < y$ if his valuation x lies strictly below y . Second, we guess and verify that, for those buyers who draw valuation $x \geq y$, the optimal bidding strategy is strictly increasing in x up to a threshold x^a , with $b(x) < a$ for $x \in [y, x^a)$ and $b(x) = a$ for all $x \geq x^a$.¹⁶ Lastly, one can easily establish that the optimal bidding strategy must satisfy $\lim_{x \rightarrow y^+} b(x) = y$.¹⁷

Therefore, the candidate bidding strategy $b(x)$ can be described:

$$b(x) = \begin{cases} 0 & \text{if } x < y \\ \widehat{b}(x) & \text{if } y \leq x < x^a \\ a & \text{if } x^a \leq x \end{cases} \quad (7)$$

with $\frac{d\widehat{b}(x)}{dx} > 0$ and $\widehat{b}(y) = y$. This strategy is an equilibrium if it is optimal for an individual buyer to bid according to $b(x)$, given that all other buyers bid according to $b(x)$. To characterize $\widehat{b}(x)$, consider the payoffs of an individual buyer who meets with a seller, draws valuation $x \in (y, x^a)$, and bids like a buyer who draws valuation $x' \in (y, x^a)$. The expected payoff to such a buyer can be written

$$u(x'|x) = \left[\frac{\lambda(1 - F(x^a))q_0(x')}{1 - q_0(x^a)} \right] [x - \widehat{b}(x')], \quad (8)$$

where the first term in equation (8) is the probability that $\widehat{b}(x')$ is the highest bid, conditional on the buyer having met with the seller, and the second term is the payoff from acquiring the good with a bid of $\widehat{b}(x')$. To understand the first term, recall from (5) that the probability that a buyer gets to meet a seller is

$$\frac{1 - q_0(x^a)}{\lambda(1 - F(x^a))}. \quad (9)$$

Since the probability of meeting with the seller *and* winning with a bid of $\widehat{b}(x')$ is $q_0(x')$, a simple application of Bayes' rule yields the first term in (8).¹⁸

¹⁶It should be fairly obvious that it is never optimal for a buyer to bid $b > a$.

¹⁷To see this, choose an $x' > y$ that is arbitrarily close to y . If $b(x') < y$, the offer is never accepted, yielding payoff zero to the buyer. If $b(x') > x'$, the buyer receives a negative payoff when the offer is accepted, and zero otherwise. Therefore, the optimal bid must lie in the set (y, x') , which is accepted with strictly positive probability, yielding a strictly positive expected payoff. Hence, as x' converges to y , the optimal bid $b(x')$ also must converge to y .

¹⁸There are several subtle points here. First, since $\widehat{b}(x)$ is assumed to be strictly increasing, the probability that $\widehat{b}(x') > \widehat{b}(x'')$ is simply the probability that $x' > x''$. Second, notice that the distribution of the number of other buyers to arrive is conditional on meeting with the seller. In other words, there is information in getting to meet with

The first-order condition with respect to x' evaluated at $x' = x$ yields the differential equation

$$\lambda f(x) \left[x - \widehat{b}(x) \right] - \frac{d\widehat{b}(x)}{dx} = 0. \quad (10)$$

Solving (10) under the boundary condition that $\widehat{b}(y) = y$ yields

$$\widehat{b}(x) = x - \frac{\int_y^x q_0(\tilde{x}) d\tilde{x}}{q_0(x)} < x. \quad (11)$$

Note that $\frac{d\widehat{b}(x)}{dx} = \frac{\lambda f(x)}{q_0(x)} \int_y^x q_0(\tilde{x}) d\tilde{x} > 0$ for all $x > y$, which confirms our initial conjecture that $\widehat{b}(x)$ is strictly increasing in x on the relevant domain.¹⁹ Finally, note that

$$\frac{d\widehat{b}(x)}{d\lambda} = \frac{\int_y^x [F(x) - F(\tilde{x})] q_0(\tilde{x}) d\tilde{x}}{q_0(x)} > 0,$$

so that buyers increase their bids in response to more (ex ante) expected competition from other buyers.

Substituting (11) into (8) yields the buyer's *ex post* expected payoff from meeting a seller and drawing valuation $x \in (y, x^a)$:

$$u(x) = \frac{\lambda(1 - F(x^a))}{1 - q_0(x^a)} \int_y^x q_0(\tilde{x}) d\tilde{x}. \quad (12)$$

Now, given the monotonicity of $\widehat{b}(x)$, the threshold x^a is the value of x such that the buyer is indifferent between acquiring the good with certainty at price a , or offering $\widehat{b}(x^a) < a$ and only acquiring the good when no subsequent buyers place a higher bid; that is, x^a satisfies

$$x^a - a = \frac{\lambda(1 - F(x^a))q_0(x^a)}{1 - q_0(x^a)} (x^a - \widehat{b}(x^a)). \quad (13)$$

Plugging in (11) yields a simple relationship between the asking price a and the cutoff x^a for any queue length λ :

$$a = x^a - \frac{\lambda(1 - F(x^a))}{1 - q_0(x^a)} \int_y^{x^a} q_0(x) dx. \quad (14)$$

the seller in the first place: it changes the probability distribution of n other buyers also arriving at the seller. All of this is incorporated into the buyer's optimal bidding strategy.

¹⁹Also note that $\frac{du(x'|x)}{dx'} = \frac{\lambda^2(1-F(x^a))}{1-q_0(x^a)} f(x') q_0(x') (x - x')$ is positive for $x' < x$ and negative for $x' > x$, confirming that $x' = x$ is the global maximum.

Differentiating (14) confirms that this relationship is one-to-one, since

$$\frac{da}{dx^a} = \frac{1 - q_0(x^a) - q_1(x^a)}{1 - q_0(x^a)} \left(1 + \frac{\lambda f(x^a)}{1 - q_0(x^a)} \int_y^{x^a} q_0(x) dx \right) > 0.$$

Hence, given any a and λ , the buyer's optimal bidding function is completely characterized by (7), where $\widehat{b}(x)$ is given by (11) and x^a is determined by (14).²⁰

Given this optimal bidding behavior, along with (9), we can calculate the ex ante expected utility that a buyer receives from visiting a seller who has posted an asking price a and attracts a queue length λ ,

$$U(a, \lambda) = \frac{1 - q_0(x^a)}{\lambda (1 - F(x^a))} \left[\int_y^{x^a} u(x) dF(x) + \int_{x^a}^{\bar{x}} (x - a) dF(x) - k \right], \quad (15)$$

where, in a slight abuse of notation, $x^a \equiv x^a(a, \lambda)$ is the implicit function defined in (14). The optimal search behavior of each buyer can then be described as follows: given the posted asking prices and the search behavior of other buyers, an individual buyer visits the seller (or mixes between the sellers) that maximizes $U(a, \lambda)$.

Seller's Asking Price. Given the optimal search and bidding behavior of buyers, we can now characterize the profit-maximizing asking price set by sellers. As a first step, we must derive the expected revenue of a seller with any asking price a and queue length λ . Recall that a seller who receives no buyers consumes his good and receives payoff y . Alternatively, if $n > 0$ buyers arrive, then the probability that all of them have a valuation below x is given by $F(x)^n$. Therefore, the density of the maximum valuation among n buyers is $nf(x)F(x)^{n-1}$, and the expected revenue of a seller with n buyers with cutoff x^a can be written as

$$F(y)^n y + \int_y^{x^a} \widehat{b}(x) nf(x) F(x)^{n-1} dx + (1 - F(x^a)^n) a.$$

Taking the expectation over n and simplifying yields

$$R(a, \lambda) = q_0(y) y + \lambda \int_y^{x^a} f(x) q_0(x) \widehat{b}(x) dx + (1 - q_0(x^a)) a, \quad (16)$$

where, again, $x^a \equiv x^a(a, \lambda)$ denotes the optimal cutoff for buyers characterized in (14), while $\widehat{b}(x)$

²⁰Notice that $\widehat{b}(x) < x$ for all $x \in (y, x^a)$, and $\lim_{x \rightarrow x^a} \widehat{b}(x) < a$. The former result resembles standard bidding behavior in first-price auctions. The latter result is to be expected as well: otherwise, for buyers with valuation x arbitrarily smaller than x^a , an arbitrarily small increase in their bid would yield a discrete increase in the probability of trading, and hence a discrete increase in their expected payoff.

denotes the optimal bidding function characterized in (11).

Note that the partial derivative of R with respect to a is strictly positive. This implies that if the queue length λ is constant, a seller who increases his asking price will always increase his revenue. Thus, if there is no relationship between a and λ — for example, if buyers search randomly across sellers — sellers behave like monopolists: they choose a sufficiently high asking price a so that $x^a = \bar{x}$. In this case, the seller meets with all buyers with probability one before choosing a trading partner, i.e., the seller runs a standard (first-price) auction.

However, in our setup, there is competition amongst sellers, who thus need to take into account that their choice of the asking price a will affect the expected number of buyers that will visit them, λ . To understand the relationship between a and λ , let us denote by \bar{U} the highest level of utility that buyers can obtain in this market, given the asking prices posted by all other sellers; as is common in this literature, we will refer to \bar{U} as the “market utility.” Then, for any $\lambda > 0$, the relationship between a and λ will be determined by the equality

$$U(a, \lambda) = \bar{U}. \quad (17)$$

In words, buyers will adjust their search behavior in such a way to make themselves indifferent between any seller that they visit with positive probability. Hence, the implicit function defined in (17) 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).

Hence, one can reinterpret the seller’s problem as a choice over both the asking price and the queue length in order to maximize revenue, subject to (17), which requires that the combination of a and λ provides the buyers with a payoff of \bar{U} . The corresponding Lagrangian can be written

$$\mathcal{L}(a, \lambda, \mu) = R(a, \lambda) + \mu [U(a, \lambda) - \bar{U}]. \quad (18)$$

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 (18), 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 (18), 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*

$$a = a^* \equiv x^* - \frac{\lambda(1 - F(x^*))}{1 - q_0(x^*)} \int_y^{x^*} q_0(x) dx, \quad (19)$$

$x^a = x^*$ 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. Proposition 2 establishes that the equilibrium implements the solution to the planner's problem. As we discussed in the introduction, this result is somewhat surprising since sellers have access to only a simple, uni-dimensional instrument, while constrained efficiency imposes requirements on two distinct margins. Before proceeding, we offer a brief discussion on the role of the asking price, and highlight a few features of our environment that allow this pricing mechanism to achieve the constrained efficient allocation in equilibrium.

First, the fact that sellers compete for buyers implies that the number of buyers who arrive at a particular seller, or queue length, depends on the seller's asking price. In typical models of directed search, in which the size of the surplus that exists between a seller and a given number of buyers is independent of the seller's pricing scheme, this relationship is straightforward. However, in our setting, sellers have to consider how their choice of a affects both the share of the surplus that buyers expect to receive *and* the size of the surplus itself (through the effect on the cutoff, x^a).

One way to understand these two channels more clearly is to use the relationship

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

to rewrite the Lagrangian (18), where $x^a(a, \lambda)$ is defined in (14). Taking first order conditions and substituting the constraint, (17), we get that profit maximization requires

$$\frac{\partial \lambda}{\partial a} \left[\frac{\partial S}{\partial \lambda} - \bar{U} \right] + \frac{\partial S}{\partial x^a} \frac{dx^a}{da} = 0. \quad (20)$$

For the sake of illustration, suppose the asking price a is greater than a^* , so that $x^a > x^*$, and consider the effects of a marginal decrease in a . First, the queue length λ increases, so that demand rises. Holding x^a fixed, the marginal effect on the seller's expected profit depends on the elasticity of this demand, $\frac{\partial \lambda}{\partial a}$, and the additional profit earned (or lost) from a marginal increase in queue length; each additional buyer that the seller is able to attract "creates" $\frac{\partial S}{\partial \lambda}$ additional revenue but "costs" \bar{U} , the market utility. Second, x^a falls, which decreases the probability that a buyer incurs the cost k when he has relatively little chance of receiving the good. Hence, holding λ fixed, lowering the asking price closer to a^* increases the size of the expected surplus. As the seller is the residual claimant on all additional surplus he creates, this increases the seller's expected profits one-for-one.

In addition to illustrating how the choice of a affects a seller's expected profits, equation (20) also lays bare the heavy burden required to achieve the constrained efficient allocation. In particular, 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 equates to ensuring that $\frac{\partial S}{\partial x^a} = 0$, so that the asking price implements the stopping rule $x^a = x^*$. The latter demands that, in equilibrium, $\frac{\partial S}{\partial \lambda} = \bar{U}$; this is a standard condition for the ex ante efficient allocation of buyers across sellers, as it implies that the buyer receive (in expectation) his marginal contribution to the match.²¹ A priori, there is no reason to think that one value of a will ensure that both of these conditions are satisfied.²² Clearly, the relationship between a and x^a has to be just right in order to ensure that both conditions are satisfied with a single choice of a .

A quick glance at equation (14) reveals that the *information structure* is crucial in determining this relationship between a and x^a . In particular, this relationship depends on the probability that a buyer wins the good given only the information contained in the fact that he has met with the seller. 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, buyers 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, as they show, each buyer has an incentive to place a bid that inefficiently pre-empts further participation. Therefore, if buyers have *too much information* then too few buyers ultimately inspect the good.

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.²³ However, he is not able to infer more, which eliminates the strategic motive of previous bidders to pre-empt entry. 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. Moreover, the size of the jump bid is such that each buyer, from an ex ante point of view, receives exactly his expected contribution to the total surplus, which ensures that the socially optimal number of buyers visit each seller to begin with.²⁴

²¹ See, e.g., Mortensen (1982), Hosios (1990), and Moen (1997).

²² The fact that $\frac{\partial S(x^*, \lambda)}{\partial \lambda} = \bar{U}$ means that the value of a inducing $x^a = x^*$ is the unique asking price that ensures each buyer in the queue receives exactly his marginal contribution to the surplus at the seller. A similar condition for efficient entry of buyers emerges in a standard private value auction without a reserve price; the intuition follows from the fact that the payoff of the winning buyer in second-price (or any revenue-equivalent) auction equals exactly his contribution to the surplus. However, we stress that for the case of a mechanism with strictly positive inspection cost, where inspection occurs sequentially and decisions are conditional on the realized outcomes of previous inspections, no such result has been proven before.

²³ In a related environment with two bidders, Pans (2013) also finds that partial disclosure is optimal for the seller.

²⁴ To see that the information structure alone is insufficient to guarantee that a simple asking price is the optimal mechanism, see Milgrom (2004, page 233). He considers an environment in which buyers do not observe their place in the queue (as in our model), but the number of potential buyers is exogenous. One can show that, in this environment, a seller who has access to only a simple asking price mechanism will again face a trade-off between

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 important results. First, even when sellers are free to post arbitrary mechanisms, all sellers posting the optimal asking price mechanism 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, we show that *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 in the real world suggests that they are a robust and compelling way to deal with the frictions in our environment.

Mechanisms. A mechanism m specifies an extensive form game that determines how a seller will select a trading partner among the buyers that visit him, as well as the payoffs that the seller and each of the buyers will receive. We allow a seller’s mechanism to condition the allocation and the payoffs on the number of buyers to arrive, along with their actions, but not on the buyers’ identities; identical buyers must be treated symmetrically.²⁵

Let \mathcal{M} be the set of all feasible mechanisms. This set includes the equilibrium asking price mechanism, which we denote by m_a^* , other asking price mechanisms (i.e., with $a \neq a^*$), as well as completely different mechanisms.²⁶ We denote the ex ante expected payoff of a seller who posts a mechanism $m \in \mathcal{M}$ and attracts a queue λ by $\mathbf{R}(m, \lambda)$, while $\mathbf{U}(m, \lambda)$ represents the expected payoff of each buyer in his queue.²⁷ The total payoffs (net of y) cannot exceed the amount of

revenue maximization and efficiency, as in Bulow and Klemperer (2009); specifically, in Milgroms environment, an unrestricted seller would optimally impose meeting fees on each buyer who inspected the good. Competition among sellers, as in our environment, eliminates these fees.

²⁵As common in the literature on competing mechanisms (see, e.g., Eeckhout and Kircher, 2010), we also require mechanisms to be anonymous in the sense that they do not condition the allocation and payoffs on the other mechanisms being posted.

²⁶For example, mechanisms including a reservation price or a participation fee.

²⁷Following Eeckhout and Kircher (2010), we assume that the seller can post an equilibrium selection device if the game has multiple equilibria, such that the payoffs are uniquely determined. In general, multiplicity is not a major

surplus $\mathbf{S}(m, \lambda)$ generated by the allocation implemented by the mechanism, which in turn cannot exceed the surplus created by the planner's solution, $\mathbf{S}^*(\lambda) \equiv S(x^*, \lambda)$. That is,

$$\mathbf{R}(m, \lambda) + \lambda \mathbf{U}(m, \lambda) - y \leq \mathbf{S}(m, \lambda) \leq \mathbf{S}^*(\lambda) \quad (21)$$

for all m and λ .

Clearly, Pareto optimality requires that the entire surplus be divided between the seller and the buyers, so we restrict attention to mechanisms that satisfy the first condition in (21) with equality. If the second condition in (21) also holds with equality, i.e., the mechanism creates the same amount of surplus as the planner's solution, we will call the mechanism *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 equilibrium strategy if, and only if, it creates the same surplus and the same expected payoffs as the optimal asking price mechanism m_a^* .

Proposition 3. *A mechanism $m \in \mathcal{M}$ is an equilibrium strategy if, and only if, it is payoff-equivalent to the optimal asking price mechanism, i.e., $\mathbf{R}(m, \lambda) = \mathbf{R}(m_a^*, \lambda)$ and $\mathbf{U}(m, \lambda) = \mathbf{U}(m_a^*, \lambda)$.*

Since the optimal asking price mechanism is surplus-maximizing, the following result is immediate.

Corollary 1. *Any equilibrium mechanism $m \in \mathcal{M}$ is surplus-maximizing.*

The intuition behind Proposition 3 and Corollary 1 is illustrated in Figure 2. 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 $\mathbf{R}(m_1, \lambda_1) - y = \mathbf{S}(m_1, \lambda_1) - \lambda_1 \bar{U}$.

INSERT FIGURE 2 HERE

The first result is that m_1 must be surplus-maximizing. To see why, suppose it is not; that is, suppose $\mathbf{S}(m_1, \lambda_1) < \mathbf{S}^*(\lambda_1)$, corresponding to point 1 in Figure 2. This seller's expected profits correspond to the intersection of the vertical axis with the line through point 1 that has slope \bar{U} . One can see immediately that this cannot be consistent with equilibrium behavior. For example,

concern, since there always exist other mechanisms that implement the desired payoffs uniquely.

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 in the figure. Since this deviation increases the size of the surplus while holding constant the market utility received by the buyers, 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 that²⁸

$$\left. \frac{d\mathbf{S}^*(\lambda)}{d\lambda} \right|_{\lambda=\lambda_3} = \bar{U}. \quad (22)$$

Equation (22) 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. Note that an asking price mechanism that implements x^* satisfies this condition, as we have demonstrated in the previous section. Hence, we obtain the following result.

Corollary 2. *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 $\bar{U} = \bar{U}^*$.

In equilibrium, all sellers will choose to post a surplus-maximizing mechanism, and each seller will attract the same queue length. Since the ratio of buyers to sellers is $\theta_b/\theta_s = \Lambda$, the total surplus at each seller must be $\mathbf{S}^*(\Lambda) = S(x^*, \Lambda)$ and the market utility must be $\mathbf{U}(m_a^*, \Lambda) = U(a^*, \Lambda)$. It therefore follows that the revenue at each seller must be equal to the revenue earned in the equilibrium with optimal asking prices,

$$\mathbf{R}(m_a^*, \Lambda) = R(a^*, \Lambda) = S(x^*, \Lambda) - \Lambda U(a^*, \Lambda) + y.$$

Again, the following corollary is an immediate consequence of Proposition 3.

Corollary 3. *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 (perhaps more complicated) mechanisms could also be utilized in equilibrium, but Proposition 3 implies that

²⁸Recall that $\mathbf{S}^*(\lambda) = S(x^*, \lambda)$ is strictly concave in λ .

any such mechanism 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^* that depends explicitly on the realization of buyers’ valuations. Moreover, the mechanism must also allocate the good to the agent with the highest valuation in the event that no agent draws a valuation $x \geq x^*$.

Taken together, we see that any mechanism must have sequential meetings (ruling out any form of static auction with simultaneous bidding) with *some* sort of device to elicit buyers’ private valuations, but this device must be flexible enough to allow the seller to return to each buyer should no other buyer reveal a valuation greater than x^* (ruling out any form of take-it-or-leave-it offers). These requirements place a fairly heavy burden on the mechanism’s design; the fact that this allocation can be achieved in a way that is both simple and familiar – with each seller posting a one-dimensional object a – suggests that asking prices are a natural way for sellers to deal with the frictions in this environment.

Finally, 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 4. *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 Positive Implications

The theory developed above provides a compelling framework for studying markets in which asking prices are prevalent. For one, the role of asking prices has explicit micro-foundations: by considering an environment in which buyers strategically choose which seller to visit, our framework produces an endogenous relationship between the asking prices set by sellers, the number of buyers who arrive at each seller, their subsequent bidding behavior, and ultimately equilibrium prices and allocations. Moreover, as we showed in the previous section, this type of pricing scheme is not imposed exogenously, but rather it emerges as the optimal mechanism; that is, sellers *choose* to use asking prices as an optimal response to the frictions in the market. Finally, despite the

relatively complex relationship between sellers' asking prices, buyers' search and bidding behavior, and the corresponding expected payoffs, the equilibrium is surprisingly simple and tractable. This tractability suggests that our framework could easily be incorporated into larger models (e.g., macroeconomic models that include a housing sector) or extended to allow for various types of heterogeneity (e.g., differentiated goods or preferences). Hence, we believe this model could be a useful benchmark for both theoretical and empirical work that focuses on markets in which asking prices are commonly used.

In this section, we flesh out some of the model's implications for a variety of observable outcomes, including the level of asking prices set by the sellers, the number of buyers who inspect the good at each seller, and the transaction price that is ultimately paid by the buyer who acquires the good. We study how these variables change with features of the economic 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.

Prices and Allocations. Figure 3 below plots a typical CDF of transaction prices, where we set $b = 0$ to represent sellers that do not trade.²⁹ 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 $q_0(x) = e^{-\Lambda}$, or because $n \geq 1$ buyers arrive but their valuations do not exceed the seller's valuation y , which occurs with probability $q_0(y) - q_0(x)$. 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 (11), 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 3 HERE

As we discussed above, 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 would provide a discrete increase in the probability of trade, at the cost of an arbitrarily small increase in the terms of trade.

Comparative Statics. Figures 4 and 5 below illustrate how equilibrium prices are affected by changes in the ratio of buyers to sellers (Λ) and changes in the inspection cost (k). 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. Though x^* is independent of Λ , notice that a^* is not; it is easy

²⁹For 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.

to verify that a^* is an increasing function of Λ . More generally, the distribution of prices under a larger Λ first-order stochastically dominates the distribution of prices under a smaller Λ . 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).

Alternatively, as the inspection cost k decreases, both the asking price a^* and the cut-off x^* 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.

INSERT FIGURES 4 and 5 HERE.

Information and Uncertainty. One might also be interested in the relationship between the ex ante dispersion in buyers' valuations, asking prices, and the ultimate transaction prices. For example, suppose 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 this information, suppose now that a buyer can identify a fraction η of goods that he would prefer, and a fraction $1 - \eta$ of goods that he would not. Formally, a good that is preferred will yield a valuation $x \in [\underline{x}', \bar{x}]$, while a good that is not will yield the buyer a valuation $x \in [\underline{x}, \underline{x}']$, where $\underline{x} < \underline{x}' \leq y < \bar{x}$. Whether or not a good is preferred is again i.i.d. across goods for each buyer.

One can easily show that buyers only visit preferred sellers and that the resulting equilibrium is very similar to what we characterize in Proposition 2, with the exception that the distribution F with support $[\underline{x}, \bar{x}]$ is replaced by the truncated distribution F' with support $[\underline{x}', \bar{x}]$. Figure 6 below illustrates the main effects of this technological improvement.

INSERT FIGURE 6 HERE

First, notice that fewer sellers don't trade, even though equilibrium queue lengths remain $\lambda = \Lambda$ at each seller. Intuitively, though the fraction of sellers who are visited by zero buyers remains constant, the truncated distribution implies that it is less likely for a seller to be visited by $n \geq 1$ buyers who all have a valuation strictly less than y . Second, notice that prices increase. This occurs 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, causing further increases in transaction prices. Given these two facts, clearly sellers' profits increase.

Prices and the Number of Inspections. In markets where asking prices are prevalent, the relationship between the sales price and the amount of time a good has been on the market has garnered substantial attention. Frequently, time itself is not really the object of interest, but rather it serves as a proxy for the number of meetings that occurred between a seller and potential buyers. Our model has direct, testable implications for this relationship between the sales price and the number of buyers who have inspected the good. Moreover, to the extent that each inspection might take a certain amount of time, our results could also be informative about the more indirect relationship between sales price and time-on-the-market, which could be helpful when the number of buyers who inspect the good is not directly observable.

Let $p(i)$ denote the expected price of a good that was sold after i meetings. There are two possible scenarios under which this event could occur. First, it could be the case that $n \geq i$ buyers arrived at this seller, the first $i - 1$ buyers had a valuation strictly less than x^* , and the i^{th} buyer had valuation $x \geq x^*$. This occurs with probability

$$\left[1 - \sum_{j=0}^{i-1} \frac{e^{-\lambda} \lambda^j}{j!} \right] F(x^*)^{i-1} [1 - F(x^*)]. \quad (23)$$

Alternatively, it could be the case that exactly $n = i$ buyers arrive, all of whom had valuation strictly less than x^* , but at least one of whom had valuation $x > y$. This occurs with probability

$$\frac{e^{-\lambda} \lambda^i}{i!} [F(x^*)^i - F(y)^i]. \quad (24)$$

In this case, the winning bid is given by $\widehat{b}(x)$, where x is drawn from the conditional distribution

$$\frac{iF(x)^{i-1}f(x)}{F(x^*)^i - F(y)^i}.$$

Therefore, the expected price of a good sold after i meetings is

$$p(i) = \frac{\left[1 - \sum_{j=0}^{i-1} \frac{e^{-\lambda} \lambda^j}{j!} \right] F(x^*)^{i-1} [1 - F(x^*)] a^* + \frac{e^{-\lambda} \lambda^i}{i!} \int_y^{x^*} iF(x)^{i-1} f(x) \widehat{b}(x) dx}{\left[1 - \sum_{j=0}^{i-1} \frac{e^{-\lambda} \lambda^j}{j!} \right] F(x^*)^{i-1} [1 - F(x^*)] + \frac{e^{-\lambda} \lambda^i}{i!} [F(x^*)^i - F(y)^i]}. \quad (25)$$

The relationship between the sales price and the number of meetings depends on the strength of two opposing forces. On the one hand, conditional on a seller not trading at the asking price, the maximum expected offer is *increasing* in i . This is a fairly obvious result: conditional on

all buyers' valuations being strictly less than x^* , the distribution of the maximum offer amongst i buyers first order stochastically dominates the distribution of the maximum offer amongst i' buyers for any $i' < i$. This effect implies that p should be increasing in i . However, on the other hand, the probability that a seller does trade at the asking price, conditional on trading after i meetings, is *decreasing* in i .³⁰ In words, as i gets larger, a trade that occurs after i meetings is more likely to be the result of the seller exhausting the queue of buyers (and trading at a price less than a), as opposed to a trade that occurred because the i^{th} buyer paid the asking price. This second effect implies that p should be decreasing in i .

The strength of each of these effects varies across different values of i , which implies an interesting shape for $p(i)$; figure 7 below provides a typical example. For small values of i , the latter effect discussed above dominates; the conditional probability of trading at price a drops off relatively quickly in the first few meetings. However, as i increases, this probability decreases at a much slower rate, and the positive effect of receiving more bids comes to dominate, causing the expected transaction price to ultimately rise in i .

INSERT FIGURE 7 HERE

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 we believe this prediction to be adequate for certain markets, it is perhaps less desirable for the analysis of other markets.³¹ It is therefore important to emphasize that our model can easily explain such transaction prices.³² To see this, note 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

³⁰Note that this probability is easy to calculate, given the expressions in (23) and (24).

³¹For example, Case and Shiller (2003) and Han and Strange (2013) provide evidence for a variety of cities in which houses occasionally sell above the asking price.

³²We would like to thank Rob Shimer for pointing out this feature of our model.

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.

While this mechanism yields the same *expected* payoffs as our (first-price) asking price mechanism, the distributions of *realized* transaction prices will differ. Figure 8 shows a typical CDF of transaction prices, where we again 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). The probability of this event is strictly positive because $a < x^a$. The following lemma presents an explicit expression.

Lemma 3. *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*

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

In other words, this interpretation 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.

INSERT FIGURE 8 HERE

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 exist institutions that make it costly to renege on an asking price. For example, as Stacey (2012) points out, real estate agents in the housing market can serve as commitment devices. For one, sellers are typically required to pay their real estate agent a commission if they receive a *bona fide* offer at the asking price, whether or not the offer is accepted. Even without this contractual clause, reputation concerns can help enforce commitment: if a homeowner repeatedly turns down buyers who offer the posted asking price, he will not only discourage brokers from listing his property, but he may also discourage other buyers from making offers in the future.

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. Of course, there are other markets in which sellers may not be able to commit to such a device; extending our framework to allow for limited commitment, or even no commitment (as in Kim and Kircher, 2012), could potentially allow our model to capture these features as well. This extension is left for future work.³³

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 becomes too large, there may exist states of the world in 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. That is, under the condition below, inspection before submitting a bid or trading is optimal both for the planner and for the buyers in the decentralized market.

Lemma 4. *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*

$$k < \int_x^y (y - x) f(x) dx. \quad (26)$$

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 (26). In words, this inequality implies that inspection is always optimal as long as 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. Before proceeding, we highlight two important points regarding (26). First, inspection will, in general, almost always remain an optimal strategy for buyers in large regions of the parameter space that do not satisfy this condition. Second, for most of the goods we have in mind (e.g., a house or a car), the assumption that the inspection cost k is small relative to the potential gains (or losses) from trade seems to be appropriate.

Free Entry. A third assumption that we discuss is our restriction to a fixed ratio of buyers to sellers. Though Proposition 2 establishes that the equilibrium cutoff x^a is efficient in our environment, a

³³A stylized but particularly tractable approach would be to assume that, after the arrival of the buyers, the seller can renege on the asking price at a certain cost or in certain states of the world. This would not only allow our model to accommodate transaction prices above the asking price, but it would also create a natural intermediate case between our framework and the model by Albrecht et al. (2012a).

traditional concern in much of the search literature is whether efficiency is also achieved when the ratio of buyers to sellers is determined endogenously.

To address this issue, suppose that sellers can freely enter the market by paying a cost c , as in standard search models (see, e.g., Pissarides, 1985). The planner will then choose Λ to maximize net social surplus, $\max_{\Lambda \in (0, \infty)} \frac{\theta_b}{\Lambda} [S(x^*, \Lambda) - c]$, while the equilibrium buyer/seller ratio follows from the indifference condition $R(a^*, \Lambda) - c = y$. In the appendix, we show that the market tightness achieved in equilibrium indeed coincides with the solution to the planner’s problem.

Lemma 5. *The market equilibrium with free entry is constrained efficient.*

Heterogeneous Inspection Costs. We chose to endogenize the number of buyers at each seller by studying a model of directed search, where a large number of sellers compete for buyers, and buyers choose strategically among these sellers. In this setting, a buyer ultimately pays two costs before acquiring a good: the inspection cost, which we discussed at length above, as well as an opportunity cost of visiting the seller (instead of a different seller), which is the equilibrium market utility \bar{U} . An alternative approach would have been to consider a single seller and replace the endogenous opportunity cost with an exogenous entry cost. Within our benchmark model, this alternative would have yielded essentially the same results: a simple asking price would have ensured that the socially optimal number of buyers would have visited the monopolistic seller, and that the socially optimal number of buyers would have inspected the good.

However, if we relax the assumption that buyers are homogeneous, then casting our environment within a model of directed search provides a distinct advantage, as the possibility of endogenous market segmentation emerges. This could help to explain why identical sellers would post different asking prices, or even use different mechanisms altogether. As a result, some goods would sell “quickly” (i.e., with high probability) at lower prices, while other goods would sell slowly at higher prices.

For example, suppose buyers are heterogeneous with respect to their inspection costs: a fraction $\phi \in (0, 1)$ of buyers have inspection cost k_1 while the remaining buyers have inspection cost k_2 , with $k_2 > k_1 > 0$. Then, as the following lemma reports, a separating equilibrium arises in which buyers sort themselves into different submarkets on the basis of their cost. In one submarket, sellers post low asking prices, which attract the buyers with high inspection costs, since they place a premium on reducing the risk of a fruitless inspection. In the other submarket, sellers post higher asking prices and serve the buyers with low inspection costs; in the limit, as these inspection costs converge to zero, the mechanism at these sellers converges to a standard auction.³⁴ In order to

³⁴In this sense, our work is related to Eeckhout and Kircher (2010). By studying various exogenous meeting technologies, they analyze how the degree of crowding out (‘rivalry’) in meetings between buyers and sellers affects

ensure that sellers are indifferent between these two mechanisms, the queue length adjusts: the sellers posting low asking prices attract relatively longer queue lengths and sell quickly.

Lemma 6. *In a market with two types of buyers, with $0 < k_1 < k_2$, there exists an asking price equilibrium with two submarkets. In each submarket $i \in \{1, 2\}$, sellers post asking price a_i and attract a queue length λ_i that consists of only buyers with inspection cost k_i , which implements a stopping rule x_i^a . This equilibrium is efficient and has the following properties: $a_1 > a_2$, $x_1^a > x_2^a$, and $\lambda_1 < \lambda_2$.*

This simple example highlights one dimension along which the model could be easily extended to confront empirical regularities in the data, such as heterogeneous asking prices or time-on-the-market, along with the coexistence of identical goods being sold with different pricing mechanisms.³⁵

8 Conclusion

In the majority of existing economic models, it is assumed that either (i) non-negotiable prices are set by sellers; (ii) prices are the outcome of a bargaining game between a single buyer and seller; 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. Without a coherent theory, it's hard for economists to identify the underlying cause for observed changes in prices and allocations in these markets, or to forecast the effects of future shocks. For example, in real estate markets, the ratio of sales prices to list prices varies widely across locations and can often change significantly over time; the three standard pricing mechanisms described above simply offer no way to interpret data of this sort.

the optimal sales mechanism. While crowding out always reduces welfare in their environment, we show that it may arise endogenously as the optimal response to the existence of inspection costs. Although our environment is quite different from theirs, we also find that the optimal sales mechanism is sometimes an auction and sometimes a posted price. In addition, we show that the asking price mechanism convexifies the range between these two extremes.

³⁵Other forms of heterogeneity, e.g. observable differences in sellers' own valuations, will often also lead to a separating equilibrium in which each submarket looks exactly as derived in section 4.

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. Given these features, we believe that our model could be useful for both theoretical and empirical work that focuses on markets in which asking prices are commonly used.

Appendix

Proof of Lemma 1. In general, if n buyers arrive at a seller, the planner might want to first learn the valuations of $i_1 \leq n$ buyers simultaneously at cost $i_1 k$. Then, depending on the realization (x_1, \dots, x_{i_1}) , the planner could choose to either allocate the good to one of these i_1 buyers (or the seller), or alternatively the planner could choose to continue and sample $i_2 \leq n - i_1$ buyers, in which case this iterative process would continue. This is precisely a special case of the problem that Morgan and Manning (1985) study, so we sketch the intuition for our result below and refer the reader to their paper for a more rigorous treatment.

To see that it can never be optimal to learn the valuation of more than one buyer at a time, suppose n 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 $\hat{x}_i \equiv \max\{y, x_1, \dots, x_i\}$.

We will show that, for *any* realizations of the first two buyers' valuations (x_1, x_2) , the net social surplus is weakly larger from learning these two valuations sequentially than it is from learning them simultaneously. To see this, realize that the net social surplus from learning these valuations simultaneously is

$$-2k - y + \max\{\hat{x}_2, Z_2(\hat{x}_2)\}, \quad (27)$$

whereas the surplus from learning these valuations sequentially is

$$\begin{aligned} & -k - y + \max\{\hat{x}_1, -k + \max\{\hat{x}_2, Z_2(\hat{x}_2)\}\} \\ = & -2k - y + \max\{\hat{x}_1 + k, \hat{x}_2, Z_2(\hat{x}_2)\} \end{aligned} \quad (28)$$

Clearly, the expression in (28) is weakly larger than (27) 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 $m > 2$ buyers simultaneously after any number $i \leq n - m$ buyers have already inspected the good. ■

Proof of Lemma 2. The proof of Lemma 2 utilizes 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 $V_{n,n}(\hat{x}_n) = \hat{x}_n - y$. As in the text, to economize on notation we will drop the second subscript “ n ” below, so that $V_{i,n}(\hat{x}_i) \equiv V_i(\hat{x}_i)$, $x_{i,n}^p \equiv x_i^p$, and so on.

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

$$Z_{n-1}(\hat{x}) = -k + \hat{x}F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} xf(x) dx. \quad (29)$$

Notice immediately that $Z_{n-1}(x^*) = x^*$ and $\frac{dZ_{n-1}(\hat{x})}{d\hat{x}} = F(\hat{x}) \in (0, 1)$, so that clearly $x_{n-1}^p = x^*$ 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 $\frac{dZ_{n-2}(\hat{x})}{d\hat{x}} = \frac{dZ_{n-1}(\hat{x})}{d\hat{x}} = F(\hat{x}) \in (0, 1)$ 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 $x_{n-2}^p = x^*$ is optimal.

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

$$Z_{n-2}(\hat{x}) - y = -k + V_{n-1}(\hat{x})F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} V_{n-1}(x)f(x) dx. \quad (30)$$

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

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

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

We have established that the following is true for $j' = 2$:

1. $Z_{n-j'}(\hat{x}) = Z_{n-j'+1}(\hat{x})$ for all $\hat{x} \geq x^*$.
2. $Z_{n-j'}(\hat{x}) > \hat{x}$ for all $\hat{x} < x^*$.
3. $\lim_{\hat{x} \rightarrow x^*} Z_{n-j'}(\hat{x}) = Z_{n-j'+1}(x^*) = x^*$,

so 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

$$Z_{n-j-1}(\hat{x}) - y = -k + V_{n-j}(\hat{x}) F(\hat{x}) + \int_{\hat{x}}^{\bar{x}} V_{n-j}(x) f(x) dx. \quad (31)$$

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

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

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 $x_{n-j-1}^p = x^*$. ■

Proof of Proposition 1. Given the results in Lemmas 1 and 2, all that remains to be shown is that it is optimal to assign queue lengths $\lambda = \Lambda$ at each seller. 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

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

Hence, the surplus generated by a seller is strictly concave in the queue length. Total surplus is therefore maximized by assigning the same queue length $\lambda = \Lambda$ to all sellers. ■

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. 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 $\widehat{b}(x)$ from (11) and a from (14) into $R(a, \lambda)$, as given in (16), 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:

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

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 $Q_1(x^a) < 1$. Therefore, the first-order conditions of the Lagrangian with 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 (14). ■

Proof of Proposition 3. Consider a candidate equilibrium in which one or more sellers post a particular trading mechanism m_1 and attract a queue $\lambda_1 > 0$, which yields them a payoff $\mathbf{R}(m_1, \lambda_1)$, yields the buyers a market utility $\bar{U} \equiv \mathbf{U}(m_1, \lambda_1)$, and creates a surplus $\mathbf{S}(m_1, \lambda_1) =$

$\mathbf{R}(m_1, \lambda_1) + \lambda_1 \mathbf{U}(m_1, \lambda_1) - y$. Equilibrium requires that no profitable deviation exists. Hence, a deviant seller who posts a mechanism m_3 and attracts a queue λ_3 , determined by market utility, must obtain a weakly lower payoff, i.e., $\mathbf{R}(m_3, \lambda_3) \leq \mathbf{R}(m_1, \lambda_1)$ for all m_3 and λ_3 such that $\mathbf{U}(m_3, \lambda_3) = \bar{U}$.

Using the definition of $\mathbf{S}(m_3, \lambda_3)$, the deviant's payoff equals

$$\mathbf{R}(m_3, \lambda_3) = \mathbf{S}(m_3, \lambda_3) - \lambda_3 \bar{U} + y.$$

This payoff is maximized when the deviant chooses a surplus-maximizing mechanism which is such that λ_3 satisfies the first-order condition

$$\left. \frac{d}{d\lambda} S^*(\lambda) \right|_{\lambda=\lambda_3} = \bar{U}. \quad (35)$$

As follows from previous results, an asking price that implements the cutoff x^* satisfies this condition and is therefore a solution to the deviant's maximization problem.

Note that (m_1, λ_1) can be part of an equilibrium if, and only if, it is a solution to the deviant's maximization problem, i.e., m_1 is surplus-maximizing and λ_1 solves (35). Since $S^*(\lambda)$ is strictly concave, a unique solution exists to (35) for each level of market utility. Hence, all sellers must attract the same queue length and this queue length must equal $\lambda = \Lambda$, since any other value would be inconsistent with the aggregate buyer-seller ratio. Since $\left. \frac{d}{d\lambda} S^*(\lambda) \right|_{\lambda=\Lambda} = \mathbf{U}(m_a^*, \Lambda)$, this implies that a mechanism m is an equilibrium strategy if and only if it creates the same surplus as the optimal asking price mechanism as well as the same payoff for the buyers. Equivalence of the seller's payoff then follows immediately. ■

Proof of Lemma 3. 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

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

which simplifies to

$$F(x^a)^n - n(F(x^a) - F(a))F(a)^{n-1} - F(a)^n.$$

Taking the expectation over n yields the desired expression. ■

Proof of Lemma 4. For the planner's problem, the proof proceeds by induction, much like the proof of lemma 2. 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.³⁶

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

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

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} , (26) 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 2.

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 choice between these options is equivalent to choosing a type of buyer $x' \in [\underline{x}, \bar{x}]$ to mimic.

We first establish that it is optimal for the deviant to behave like a buyer who has a valuation x' equal to the unconditional expected value of x , which we denote by $x^e = E_F[x] \equiv \int_{\underline{x}}^{\bar{x}} xf(x) dx$. To see this, consider the expected payoffs under each of the three options. First, a deviant who acts like a buyer with valuation $x' \in [\underline{x}, y]$ receives a payoff of 0. Second, if the deviant instead imitates a buyer of type $x' \in (y, x^a)$ and bids $\hat{b}(x')$, his payoff conditional on meeting with the seller is

$$u(x'|x^e) = \frac{\lambda(1 - F(x^a))q_0(x')}{1 - q_0(x^a)} [x^e - \hat{b}(x')].$$

³⁶The 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 2, this is dominated by learning the valuation of the last buyer since $\hat{x}_{n-1} < x^*$.

Since $\hat{b}(x')$ is optimal in equilibrium, it follows that the deviant should choose $x' = x^e$ and bid $\hat{b}(x^e)$. Evaluating the expected payoff from this strategy yields

$$u(x^e) = \frac{\lambda(1 - F(x^*))}{1 - q_0(x^*)} \int_y^{x^e} q_0(\tilde{x}) d\tilde{x}, \quad (36)$$

where $u(x^e)$ is increasing and convex in its argument. Finally, if the deviant mimics a buyer of type $x' \in [x^*, \bar{x}]$ and bids the asking price, he obtains a payoff

$$x^e - a^* = x^e - x^* + u(x^*). \quad (37)$$

Comparing the three payoffs reveals that the deviant maximizes his payoff by behaving as a buyer with valuation x^e . 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 (26), 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 distribution $\tilde{F}(x)$ 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 (26), 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 $u(x)$ is an increasing function then implies that $u(x^e) < u(\tilde{x}^e)$, while the convexity of $u(x)$ implies that $u(\tilde{x}^e) < E_{\tilde{F}}[u(x)]$ by Jensen's inequality. Note, however, that $E_{\tilde{F}}[u(x)]$

exactly equals the payoff from inspection, since

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

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

To show that equation (26) 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 (26) holds.

■

Proof of Lemma 5. The equilibrium buyer-seller ratio is determined by the free entry condition $\hat{R}(x^*, \Lambda) - c = y$. Since $S(x^*, \lambda) = \hat{R}(x^*, \lambda) + \lambda \hat{U}(x^*, \lambda) - y$, this condition is equivalent to

$$S(x^*, \Lambda) - c = \Lambda \hat{U}(x^*, \Lambda). \quad (38)$$

The efficient buyer-seller ratio follows from the first-order condition of net social surplus $\frac{\theta_b}{\Lambda} [S(x^*, \Lambda) - c]$ with respect to Λ . After simplification, this yields

$$-\frac{1}{\Lambda} [S(x^*, \Lambda) - c] + \frac{\partial}{\partial \Lambda} S(x^*, \Lambda) = 0. \quad (39)$$

Since $\frac{\partial}{\partial \Lambda} S(x^*, \Lambda) = \hat{U}(x^*, \Lambda)$, it follows that the solution to (38) solves (39). ■

Proof of Lemma 6. The first part of this proof closely resembles the proof of proposition 2, although we focus on the ‘dual’ problem, i.e., the buyers’ maximization problem. Let $\hat{R}(x^a, \lambda)$ denote a seller’s revenue given a cutoff x^a and a queue λ . Similarly, let $\hat{U}_i(x^a, \lambda)$ denote the corresponding payoff of a buyer with inspection cost k_i . The Lagrangian for a buyer of type i is

$$\hat{\mathcal{L}}_i(x^a, \lambda, \psi) = \hat{U}_i(x^a, \lambda) + \psi \left(\hat{R}(x^a, \lambda) - \bar{R} \right),$$

where ψ is the multiplier on the constraint requiring that sellers only offer combinations of (x^a, λ) that provide them with their equilibrium payoff \bar{R} . The FOCs of this Lagrangian with respect to x^a and λ correspond to the ones in the proof of proposition 2 for $\psi = \frac{1}{\mu}$. They therefore imply

$\int_{x^a}^{\bar{x}} (x - x^a) dF(x) = k_i$. Let x_i^a denote the cutoff that solves this equation. Then clearly $x_1^a > x_2^a$.

A pooling equilibrium in this environment would require that all sellers post the same asking price a and attract the same queue length Λ , consisting of both types of buyers. However, this asking price and queue length cannot simultaneously implement both x_1^a and x_2^a . Hence, a pooling equilibrium does not exist. A similar argument rules out partial pooling.

Instead, a separating equilibrium with two submarkets must arise, in which a fraction of sellers post an asking price a_i that attracts only buyers with inspection cost k_i , for $i \in \{1, 2\}$. The equilibrium queue length at each type of seller, λ_i , will then imply the cutoff x_i^a described above. Since a seller's revenue is strictly increasing in both his cutoff and queue length (see again the proof of proposition 2), the seller's indifference condition implies that the queue length must be shorter in the submarket with the higher cutoff, i.e., $\lambda_1 < \lambda_2$.

Having established that $x_1^a > x_2^a$ and $\lambda_1 < \lambda_2$, in equilibrium, it remains to be shown that $a_1 > a_2$. Suppose, towards a contradiction, that $a_2 \geq a_1$. The equilibrium revenue \bar{R}_2 for sellers in the submarket that implements x_2^a then equals

$$\bar{R}_2 = yq_0(y; \lambda_2) + \int_y^{x_2^a} \hat{b}(x; \lambda_2) dq_0(x; \lambda_2) + \int_{x_2^a}^{x_1^a} a_2 dq_0(x; \lambda_2) + \int_{x_1^a}^{\bar{x}} a_2 dq_0(x; \lambda_2),$$

where λ is (re-)introduced as an argument of the functions $\hat{b}(\cdot)$ and $q_0(\cdot)$ to avoid ambiguity.

Note that given a particular cutoff, revenue is strictly increasing in the queue length and the asking price. Note further that the asking price a_1 must exceed $\hat{b}(x_1^a)$. Hence,

$$\begin{aligned} \bar{R}_2 &> yq_0(y; \lambda_1) + \int_y^{x_2^a} \hat{b}(x; \lambda_1) dq_0(x; \lambda_1) + \int_{x_2^a}^{x_1^a} a_1 dq_0(x; \lambda_1) + \int_{x_1^a}^{\bar{x}} a_1 dq_0(x; \lambda_1) \\ &> yq_0(y; \lambda_1) + \int_y^{x_2^a} \hat{b}(x; \lambda_1) dq_0(x; \lambda_1) + \int_{x_2^a}^{x_1^a} \hat{b}(x; \lambda_1) dq_0(x; \lambda_2) + \int_{x_1^a}^{\bar{x}} a_1 dq_0(x; \lambda_1) \\ &= \bar{R}_1, \end{aligned}$$

where \bar{R}_1 denotes the equilibrium revenue for sellers in the submarket that implements x_1^a . This yields the desired contradiction. Hence, $a_1 > a_2$.

Conditional on separation, the efficiency of x_1^a and x_2^a immediately follows earlier results. Hence, the only thing left to prove is that the planner indeed prefers separation over pooling. If the planner pools both types of buyers, then ex post the optimal cutoff x^p is the solution to $\int_{x^p}^{\bar{x}} (x - x^p) dF(x) = \phi k_1 + (1 - \phi)k_2$. Output in that case equals $\theta_s S(x^p, \Lambda)$.

When the planner separates the buyer types into submarkets with cutoffs x_1^* and x_2^* instead, he needs to decide how many sellers to allocate to each submarket. Let η denote the fraction of sellers

assigned to the submarket with x_1^a . Optimal output then equals

$$\max_{\eta} \eta \theta_s S \left(x_1^*, \frac{\phi}{\eta} \Lambda \right) + (1 - \eta) \theta_s S \left(x_2^*, \frac{1 - \phi}{1 - \eta} \Lambda \right).$$

Separation creates more output than pooling, since

$$\begin{aligned} & \max_{\eta} \eta \theta_s S \left(x_1^*, \frac{\phi}{\eta} \Lambda \right) + (1 - \eta) \theta_s S \left(x_2^*, \frac{1 - \phi}{1 - \eta} \Lambda \right) \\ & > \max_{\eta} \eta \theta_s S \left(x^p, \frac{\phi}{\eta} \Lambda \right) + (1 - \eta) \theta_s S \left(x^p, \frac{1 - \phi}{1 - \eta} \Lambda \right) \\ & > \theta_s S(x^p, \Lambda). \end{aligned}$$

■

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. (2012a). Directed search in the housing market. mimeo.
- Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2012b). A note on Peters and Severinov, “Competition among sellers who offer auctions instead of prices”. *Journal of Economic Theory*, 147:389–392.
- Arnold, M. A. (1999). Search, bargaining and optimal asking prices. *Real Estate Economics*, 27(3):453–481.
- Bajari, P. and Hortacsu, A. (2003). The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2):329–55.
- 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 Sákovics, J. (1999). Imperfect competition in auction designs. *International Economic Review*, 40(1):231–247.
- Case, K. and Shiller, R. (2003). Is there a bubble in the housing market? *Brookings Papers on Economic Activity*, 2:299–342.
- 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. (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.

- Han, L. and Strange, W. (2013). What is the role of the asking price for a house? mimeo.
- Hosios, A. (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. (2012). Efficient cheap talk in directed search: On the non-essentiality of commitment in market games.
- Levin, D. and Smith, J. (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.
- McAfee, R. and McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, 25:699–738.
- McAfee, R. and McMillan, J. (1988). Search mechanisms. *Journal of Economic Theory*, 44:99–123.
- McAfee, R. P. (1993). Mechanism design by competing sellers. *Econometrica*, 61(6):pp. 1281–1312.
- Menzio, G. (2007). A theory of partially directed search. *Journal of Political Economy*, 115:748–769.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- 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. mimeo, University of Rochester.
- Peters, M. and Severinov, S. (1997). Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75:141–179.
- Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies and real wages. *American Economic Review*, 75:676–690.
- Reynolds, S. and Wooders, J. (2009). Auctions with a buy price. *Economic Theory*, 38(1):9–39.

- Roberts, J. and Sweeting, A. (2012). When should sellers use auctions? *American Economic Review*, forthcoming.
- 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.

Figure 1: Optimal Cutoff Rule

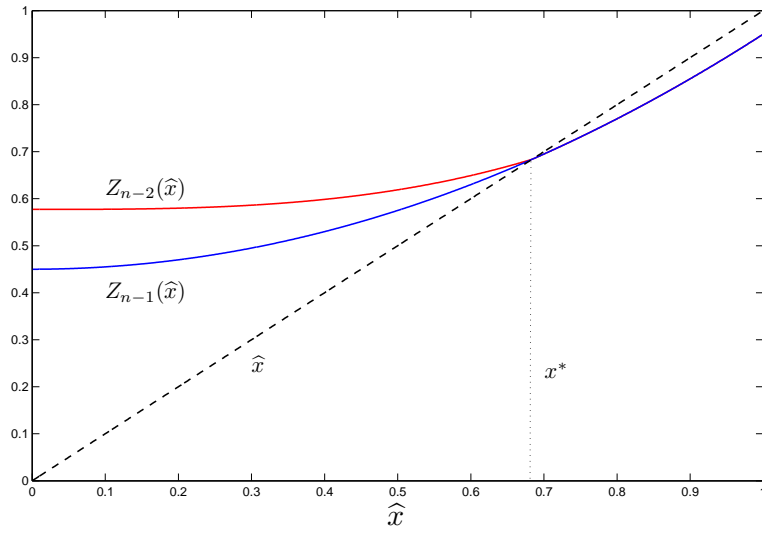


Figure 2: Optimal Mechanisms

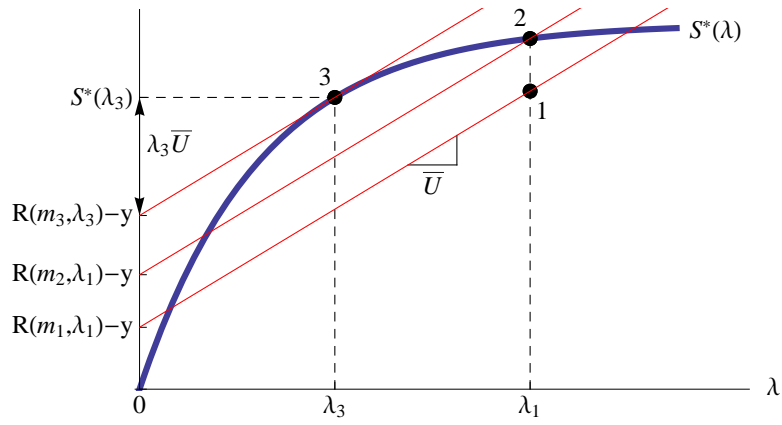


Figure 3: Transaction Price CDF

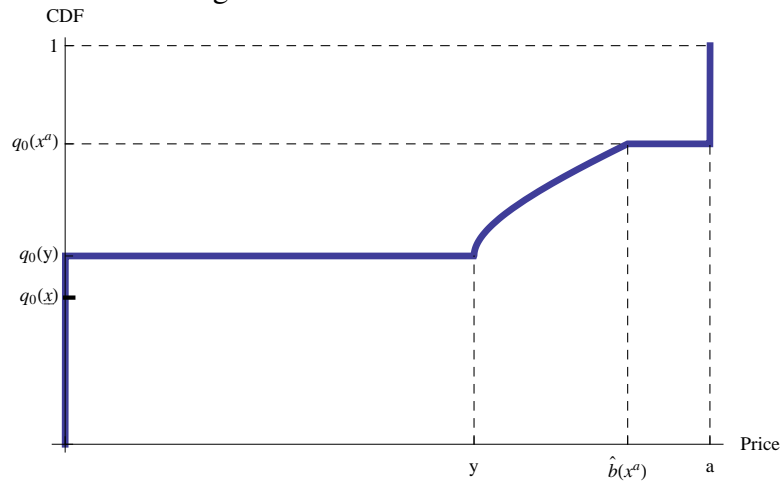
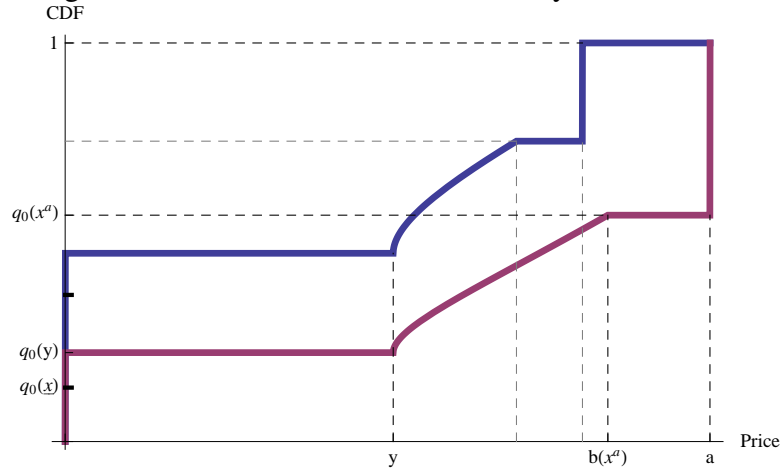
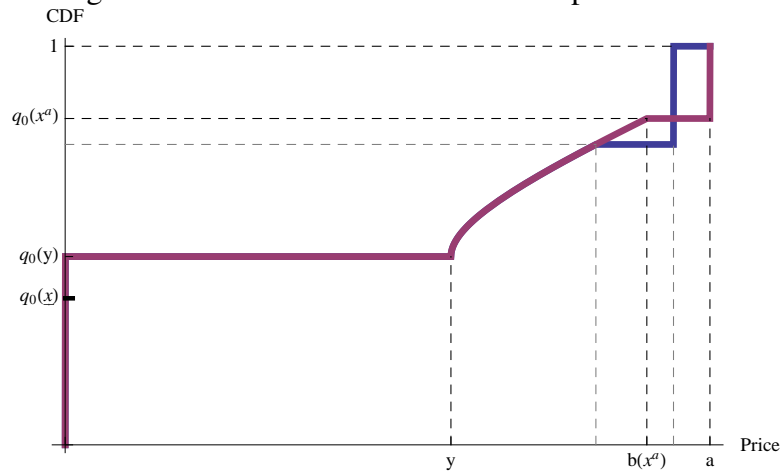


Figure 4: Effect of an Increase in the Buyer/Seller Ratio



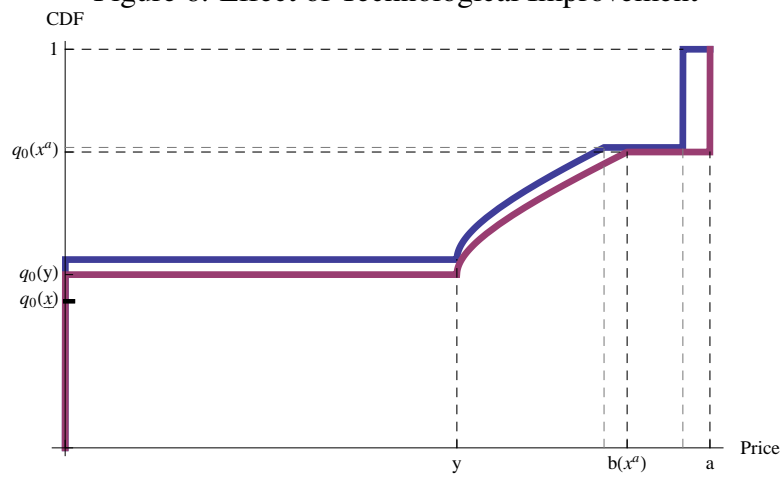
blue line = original equilibrium; purple line = equilibrium with higher Λ .

Figure 5: Effect of a Decrease in the Inspection Cost



blue line = original equilibrium; purple line = equilibrium with lower k .

Figure 6: Effect of Technological Improvement



blue line = original equilibrium; purple line = equilibrium with more information.

Figure 7: Number of Meetings and Transaction Price

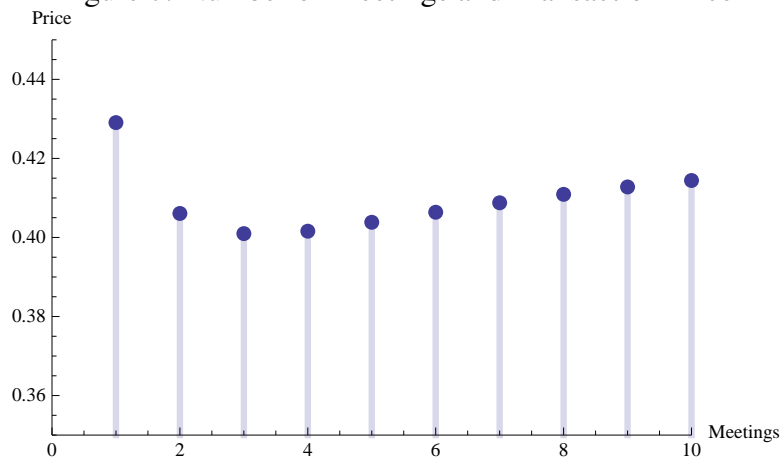


Figure 8: Transaction Prices Above the Asking Price

