Market Institutions: An Expository Essay

Singh, Nirvikar

University of California, Santa Cruz

22 March 2010

Market Institutions: An Expository Essay

Nirvikar Singh
Department of Economics
University of California, Santa Cruz

March 22, 2010

Abstract

This essay provides an elementary, unified introduction to the models of market institutions that go beyond the competitive model of price-taking behavior on both sides of the market. Several models of market institutions that govern price determination are explored and compared, including contracting, posted prices, bilateral bargaining, middlemen, and auctions. While equilibrium models still do not capture the full possibilities for market behavior, modeling specific market institutions reduces the level of abstraction inherent in the standard competitive model.

Keywords: market institutions; contracting; posted prices; bilateral bargaining; middlemen; auctions

Introduction

The standard competitive model makes two kinds of assumptions about competitive market institutions. First, price-taking behavior is literally that – both sides of the market assume that they can buy or sell what they wish at the prevailing market price. A market equilibrium is precisely a price (or prices in multiple markets) that ensures that aggregate plans of buyers and sellers are matched, so that each market participant can realize her plans. This assumption can be relaxed, to allow a richer set of contracts to be offered in the market. This possibility is alluded to at the end of Singh (2010a), in the context of an insurance market with asymmetric information, and we develop that case here. Models of contracting, especially in the context of uncertainty or asymmetric information, constitute an extensive literature. The distinction in this essay is to consider them in the context of different market institutions for price determination. Typically, textbook treatments (e.g., Kreps, 1990; Mas-Colell et al., 1995; Varian, 1992) group analyses by methodologies, rather than the idea of explaining a diverse range of institutions – hence the distinction of this essay is in its grouping of topics.[1]

A second, more significant assumption in the competitive model is that it abstracts entirely from the process of price determination. In a market with competitive buyers and sellers, everyone is a price-taker, and no one sets prices. Of course, it is well recognized that this is an abstraction – it just happens to be a very useful one for getting to the heart of the valuation and efficiency role of market-determined prices. At the same time, much thought has gone into building models in which there is a specific institutional mechanism for price determination. In some cases, the abstraction is pushed one level deeper, but is still there. In particular, equilibrium models still do not provide detailed analysis of behavior out of equilibrium and how exactly markets adjust to reach equilibrium. Nevertheless, the models provide insight into the details of market institutions, and how those details can matter for resource allocation and efficiency. Typically, these models are based on game theoretic approaches, but use assumptions that capture or approximate the impacts of competition (e.g., large numbers of market

[1] A seminal survey of models of market organization, albeit with a very different focus, was Rothschild (1973). Singh (2010b) has some overlap with the focus of Rothschild's survey.

participants). Examples of market institutions include posted prices (prices set by sellers), bargaining or negotiation, middlemen and auctions.

Contracts

A market equilibrium with contracts will be illustrated in the context of a simple insurance market model (Rothschild and Stiglitz, 1976; Singh, 2010a). The key feature of that market is an asymmetry of information about the risk level of individuals. This asymmetry makes expanding the space of market possibilities to include contracts worth considering. In the absence of the information asymmetry, one would not need contracts more complicated than agreements to buy or sell at the going price. In fact, the overall uncertainty in the insurance market (whether an individual will have an accident or not) is not an essential feature. For example, the same contracting possibilities would arise if an employer was contracting with employees – with output certain given ability – but ability was not observed by the employer. Uncertainty could still play a role, however, in making it impossible to infer ability from output.

To review the assumptions from Singh (2010a), there are two kinds of individuals, differing only in their probabilities of suffering a loss. However, market opportunities are no longer described by the ability to purchase any number of units of insurance, *I*, at a given price *P*. Instead, an insurance opportunity is described by a contract that specifies the total premium, *R*, and the total payout, *I*. An individual can choose to purchase the contract or not. Insurance companies will offer contracts, and in a competitive equilibrium, make zero expected profits.

First consider the case of only one type of individual. She chooses the insurance contract if the following holds:

$$\pi^1 U(\omega^1 - R) + (1-\pi^1)U(\omega^1 - d - R + I) \geq \pi^1 U(\omega^1) + (1-\pi^1)U(\omega^1 - d) \qquad (1)$$

The zero expected profit condition requires that $R-(1-\pi^1)I=0$. Note that the implied price of insurance in this case is $R/I=(1-\pi^1)$, as was the case with competitive insurance in Singh (2010a). Hence the left hand side of (1) becomes

$$\begin{aligned}&\pi^1U(\omega^1-(1-\pi^1)I)+(1-\pi^1)U(\omega^1-d-(1-\pi^1)I+I)\\&=\pi^1U(\omega^1-I+\pi^1I)+(1-\pi^1)U(\omega^1-d+\pi^1I)\end{aligned} \tag{2}$$

Differentiating the last expression with respect to *I* yields

$$\begin{aligned}&-\pi^1(1-\pi^1)U'(\omega^1-I+\pi^1I)+(1-\pi^1)\pi^1U'(\omega^1-d+\pi^1I)\\&=-\pi^1(1-\pi^1)[U'(\omega^1-I+\pi^1I)-U'(\omega^1-d+\pi^1I)]\end{aligned} \tag{3}$$

By risk aversion, the marginal utility is decreasing in income and so the expression in square brackets is positive (negative) if $I>(<)d$. Hence $I=d$ maximizes the left hand side of (1), and this would be the natural contract equilibrium with competitive insurance firms, i.e., $(R^*,I^*)=((1-\pi^1)d,d)$. The individual obtains full insurance, not by choosing the amount of insurance at a market price per unit, but by choosing the insurance contract that prevails in the competitive market.

Next, we turn to the case where there are two risk levels, such that $\pi_H^1<\pi_L^1$. If insurance companies observe risk types, then there can be a competitive equilibrium with two types of contracts, $(R_i^*,I_i^*)=((1-\pi_i^1)d,d)$. Each type is restricted to the contract that is tailored for that type, and everyone gets full insurance in equilibrium. The outcome is again the same as in Singh (2010a), with full information about individuals' risk levels.

In Singh (2010a), we discussed what would happen if individual risk levels were not observable by insurance companies, so that all individuals would have to be offered the same price of insurance. In that case, there could be two types of competitive equilibria: (i) both types of individuals purchase insurance, at a price that would average the probabilities of loss across the two risk levels; (ii) only high-risk individuals would purchase insurance, at a price reflecting only their own loss probabilities, with low-risk individuals choosing to purchase no insurance. The second equilibrium arises when the extra cost of insuring high-risk individuals makes insurance too expensive for the low-risk types.

We now investigate when it is possible that there might be two types of contracts, such that the equilibrium is better in some sense than one or both of the above equilibria. The possibility of specifying the total amount of insurance purchased in a contract creates this new possibility. If only prices can be specified in the market, then any buyer of insurance would choose the lower price. However, if a lower price comes with less coverage, different buyers might have different rankings of contracts available in the marketplace.

Therefore, we are looking for two contracts, $(R_L, I_L), (R_H, I_H)$, such that low risk individuals would prefer the first contract, and high-risk individuals would prefer the second contract. This implies two inequalities must hold simultaneously:

$$\begin{aligned}
&\pi_L^1 U(\omega^1 - R_L) + (1-\pi_L^1)U(\omega^1 - d - R_L + I_L) \geq \\
&\qquad \pi_L^1 U(\omega^1 - R_H) + (1-\pi_L^1)U(\omega^1 - d - R_H + I_H) \\
&\pi_H^1 U(\omega^1 - R_H) + (1-\pi_H^1)U(\omega^1 - d - R_H + I_H) \geq \\
&\qquad \pi_H^1 U(\omega^1 - R_L) + (1-\pi_H^1)U(\omega^1 - d - R_L + I_L)
\end{aligned} \tag{4}$$

These are called the *self-selection* or *incentive compatibility constraints*, in a contract equilibrium with asymmetric information. We must also note that choosing the contracts must be better for individuals of either type than not purchasing insurance at all. This requirement gives us the following two acceptance constraints, which are the analogues of (1) for the case of two risk levels.

$$\begin{aligned}
&\pi_L^1 U(\omega^1 - R_L) + (1-\pi_L^1)U(\omega^1 - d - R_L + I_L) \geq \pi_L^1 U(\omega^1) + (1-\pi_L^1)U(\omega^1 - d) \\
&\pi_H^1 U(\omega^1 - R_H) + (1-\pi_H^1)U(\omega^1 - d - R_H + I_H) \geq \pi_H^1 U(\omega^1) + (1-\pi_H^1)U(\omega^1 - d)
\end{aligned} \tag{5}$$

Now we impose the condition that expected profits are zero for each type of contract. This is what we might expect with competition among insurance firms. Therefore, $R_i - (1-\pi_i^1)I_i = 0,\ i = L, H$. This means that the expected utility with an insurance purchase has the form $\pi_i^1 U(\omega^1 - I_i + \pi_i^1 I_i) + (1-\pi_i^1)U(\omega^1 - d + \pi_i^1 I_i), i = L, H$. It also implies that the acceptance constraints will not be binding, so we can focus on the self-selection constraints.

Intuitively, the restriction imposed by the self-selection constraints will work as follows. Suppose that the full-insurance contracts that would work with the risk levels being observable were known were offered. The low-risk individual would not be interested in $((1-\pi_H^1)d, d)$, since it offers worse terms for the insurance. On the other hand, the high risk individual would always choose the contract $((1-\pi_L^1)d, d)$ over the contract $((1-\pi_H^1)d, d)$. Hence, intuitively, the first of the self-selection constraints will not be binding, and the only restriction we are left with is

$$\begin{aligned} &\pi_H^1 U(\omega^1 - (1-\pi_H^1)d) + (1-\pi_H^1)U(\omega^1 - (1-\pi_H^1)d) \\ &= \pi_H^1 U(\omega^1 - (1-\pi_L^1)I_L) + (1-\pi_H^1)U(\omega^1 - d + \pi_L^1 I_L) \end{aligned} \tag{6}$$

This says that the high-risk type must be indifferent between the two market insurance contracts in the competitive equilibrium. The left hand side is the expected utility from the full insurance contract, and just reduces to $U(\omega^1 - (1-\pi_H^1)d)$. The right hand side is what the high-risk type would get by choosing $((1-\pi_L^1)I_L), I_L)$, the contract meant for the low-risk type. The latter still has fair odds, but the quantity will be restricted, so that $I_L < d$, i.e., the low-risk type ends up with less than full insurance. This is how the quantity restriction comes into play in the contract offerings, versus the usual competitive price-taking assumption that any amount can be bought and sold at the going market price.

The idea that there is an equilibrium set of contracts in a competitive market generalizes the usual notion of an equilibrium price or set of prices. It does not require that either side of the market sets prices, nor that there is any market power that results from price setting. In that sense, the institutions of the market remain an abstraction. One could posit, in the above model, that insurance firms decide what contracts to offer. Formally, the equilibrium notion shifts from a price/contract equilibrium (is there a set of prices/contracts such that supply equals demand?) to a Nash equilibrium in firms' choices of prices/contracts. We turn next to price-setting equilibria.

Posted Prices[2]

A posted price is a price set by a firm, at which it is willing to sell. This is by far the most common institution for retail markets. Goods are on shelves or in catalogues, with marked prices, and the seller implicitly agrees to sell whatever it has available at that price. Services are also typically made available with prices publicly announced or provided to a prospective buyer on enquiry.

As the Bertrand model shows, price setting of this form does not imply market power. Even if there are just two price-setting sellers of a homogeneous good, the outcome can be one where the good is priced at marginal cost, i.e., the competitive price. However, the equilibrium can be very sensitive to the exact nature of the posting, and when consumers learn the price. In the Bertrand model, buyers know all firms' prices, and can shift costlessly and instantly to a firm that offers a lower price. This is what makes undercutting profitable if a firm's price is above marginal cost, and what drives price to marginal cost in equilibrium.

Instead, suppose that prices are posted only in the sense that a potential buyer can observe the price when she visits or contacts the seller, the price is fixed, and it is accompanied by an offer to supply the buyer's demand at that price. Suppose also that the potential buyer incurs a cost if she does not buy after learning the offer. This may be a physical or monetary cost of visiting or contacting another seller, or a time cost of delay in

[2] The model and analysis in this section are taken from Riley and Zeckhauser (1980).

purchasing. The buyer's search is sequential, but does not follow an optimal dynamic strategy, in the sense that the buyer does not estimate a distribution of prices and update it based on observation. In this model, the equilibrium price turns out to be the monopoly price, no matter how small is the buyer's search cost. So making these changes in what buyers know and how they act without perfect knowledge pushes the equilibrium form one polar case (the competitive price) to the other (the monopoly price).

The rough argument for the result is as follows. Suppose that there are n sellers, and $N(\overline{p})$ buyers in the market with reservation price less than or equal to $\overline{p}$. With identical sellers of a homogeneous good, in equilibrium the number of buyers for each seller will be $\frac{1}{n}N(\overline{p})$. The demand function that each seller faces will be $\frac{1}{n}N(\overline{p})D(\overline{p})$, and its profit (assuming constant unit cost of c) will be $\frac{1}{n}N(\overline{p})D(\overline{p})(\overline{p}-c)$. Now if potential buyers face a cost of going somewhere else, then there should be an $\varepsilon > 0$ such that buyers will still purchase from this seller rather than going somewhere else, so that $N(\overline{p}) = N(\overline{p}+\varepsilon)$. In that case, the seller needs only to maximize $D(\overline{p})(\overline{p}-c)$. But that means choosing the monopoly price, so the equilibrium price must be the monopoly price, $p^* = p^m$.

In fact, one can vary the assumptions so that neither the competitive price nor the monopoly price is the equilibrium price, in this posted price setting. There can be a single equilibrium price between the two polar cases, or even a distribution of prices, despite there being a homogeneous good. The nature of information, of how information is acquired, and how it is disseminated can all affect the predicted outcome. In a sense, the chief lesson of these models is that the nature of a posted price equilibrium is very sensitive to the detailed structure of market institutions.

One issue we have taken for granted in the above discussion is that there is no alternative to a posted price – only its level has been in question. However, one can also ask why a potential buyer, when faced with a high price offer, cannot make a counter offer, or, more importantly, why a seller would not entertain any such counter offers. In the case of a large complex organization, the answer can be found in transaction costs relative to the price of a good. It may make sense for a car salesperson to bargain in such a manner, but not for a grocery store clerk. We treat bargaining more explicitly in the next section, but first introduce a model in which the decision to choose a posted price is endogenous – the seller optimally refuses to shade the price offer.

Consider a situation where a series of buyers enter the seller's store until a sale is made: there is just a single unit to be sold. Buyers have no costs of waiting or negotiating, and they care only about the difference between their reservation value and the price they pay. The seller maximizes her expected profit. A key assumption is that the seller is able to commit to a strategy, and that each buyer knows this. Also, the seller does not know any given buyer's reservation price, but knows the distribution of reservation prices over potential buyers. If the seller knew the buyer's reservation price, she would just wait for the buyer with the highest reservation price, ask for that price and not budge.

We now sketch the argument for why the seller will choose fixed prices in this situation. Let S denote any contingent pricing strategy of the seller. For example, S might say that if she is turned down at a price of p_1, she quotes a new price of p_2 with probability 0.5 and requests a new buyer with probability 0.5. A seller with reservation price v has some optimal response to this strategy, denoted by $b_S(v)$. The seller's strategy and buyer's response together determine an implied probability of sale and expected price, as follows:

$$\begin{aligned} H(v) &\equiv H_S(b_S(v)) \\ \overline{p}(v) &\equiv \overline{p}_S(b_S(v)) \end{aligned} \tag{7}$$

The seller uses the probability of sale function in calculating expected profit of a current customer versus a future potential buyer. The key step in the argument examines the difference between the two gains, weighted by $H(v)$. Since this difference over possible intervals of v will be either positive or negative, it turns out that the seller's optimal strategy is given by

$$H^*(v) = \begin{cases} 0, & v < v^* \\ 1, & v \geq v^* \end{cases} \tag{8}$$

Mathematically, the seller solves a control problem which has a bang-bang solution. The intuition is that if there is a greater profit to be made by selling to a buyer with a particular valuation versus waiting, then it should be done for sure. Furthermore, there is a specific cutoff of valuations below which it is always better to wait. The cutoff value v^* is the take-it-or-leave-it fixed price.

Bargaining

The last model of the previous section allowed for bargaining, but in a somewhat abstract way. Furthermore, it turned out that bargaining was not an equilibrium, even when it was possible. The idea was to show that posted prices might emerge as an endogenous choice from a wider range of pricing options. In practice, bargaining does occur, and this section explores some ways of modeling the bargaining process.

The simplest bargaining model assumes perfect information, and an intuitive bargaining protocol: the players make alternating offers.[3] One player starts with an offer. If the other accepts, the bargaining is over and the deal is concluded. If the second player rejects, he gets to make an offer, in which case player 1 can accept or reject, and so on. We also assume that there is just a single unit to be bought and sold. Suppose that the buyer's valuation is v, while the seller's is 0. A transaction price p therefore assigns surplus p to the seller and $v - p$ to the buyer. Essentially, therefore, bargaining over the transaction price reduces to bargaining over splitting the fixed surplus v. Player 1 can be the seller or buyer in this case.

---

[3] This is the famous Rubinstein (1982) bargaining model, with the treatment here, as in textbooks, being influenced heavily by Shaked and Sutton (1984).

This bargaining model also assumes that there are some costs of bargaining, in the form of time discounting. Each offer takes up a period, and the payoffs in the next period are discounted by a factor $\delta \in (0,1)$, which is the same for both buyer and seller. The period can be as small a time interval as we like, and the discount factor can therefore be as close to one as we like.

First suppose that there is a finite amount of time available for bargaining, say T periods. Then the bargaining game can be solved by backward induction. If T is odd, then player 1 makes the last offer, assuming that no agreement has been reached previously. Since there is no advantage to rejecting an offer, player 2 should accept any offer, as low as v_2 = 0. Hence the payoffs are $(\delta^{T-1}v, 0)$.

In period $T - 1$, however, player 2 gets to make the offer. Player 1 should accept any offer that provides at least $\delta^{T-1}v$, since that is what she can get by rejecting and having the chance to make a final offer in period T. This is therefore what player 2 offers in period $T - 1$, and the payoffs are $(\delta^{T-1}v, \delta^{T-2}v - \delta^{T-1}v)$. Compared to the last period, player 2 has less bargaining power because player 1 now has the option of rejecting and getting to make a final take-it-or-leave-it offer. Still, player 2 has some bargaining power from being the one to make the offer.

Continuing with this backward induction, we can see that the unique Subgame Perfect Nash Equilibrium (SPNE) of this alternating offer bargaining game is where player 1 makes an offer in period 1, it is accepted, and the payoff to player 1 is:

$$\begin{aligned} v_1^*(T) &= v[1 - \delta + \delta^2 - \ldots + \delta^{T-1}] \\ &= v\left(\frac{1+\delta^T}{1+\delta}\right) \end{aligned} \tag{9}$$

Player 2's equilibrium payoff is therefore $v_2^*(T) = v - v_1^*(T)$. If player 1 is the seller, then the equilibrium sales price in this situation is $p^* = v_1^*(T)$. One can check that player 1 gets a greater share of the surplus in this equilibrium. This reflects two advantages: she gets to make the first offer, and she would get to make the last offer, if that situation arose.

If the number of time periods approaches infinity, then the bargaining equilibrium approaches

$$\left(\frac{v}{1+\delta}, \frac{\delta v}{1+\delta}\right) \tag{10}$$

In this case, the advantage of making a last offer goes away, since there is no last period, but there is still an advantage to making the first offer, and so player 1 still gets more than half of the surplus. Again, if player 1 is the seller, the equilibrium sales price is $p^* = \frac{v}{1+\delta}$.

The case of T being even can be worked out using the analysis already performed. This is because after one period, with no agreement in period 1, the subgame is one of $T-1$ periods, where $T-1$ is odd, but player 2 is making the initial offer in this subgame. Hence, player 2 gets $\delta v_1^*(T-1)$, leaving $v - \delta v_1^*(T-1)$ for player 1. This amount is explicitly

$$\begin{aligned} & v - \delta v[1-\delta+\delta^2 - \ldots + \delta^{T-2}] \\ &= v[1-\delta+\delta^2 - \ldots - \delta^{T-1}] \\ &= v\left(\frac{1-\delta^T}{1+\delta}\right) \end{aligned} \tag{11}$$

Now player 1 gets somewhat less than in the case where she makes the last offer, since player 2 would be in a monopoly position at the last bargaining stage, if that were to be reached. Even though the last period of bargaining is never reached in equilibrium, its anticipation filters through to the equilibrium through the backward induction process.

Note that, if the number of periods approaches infinity, the solution to the case with an even number of periods is identical to that with an odd number of periods. The last-period effect goes away, leaving only the impact of the advantage of making the first offer. It also turns out that the equilibrium split of the surplus in (10) is the unique equilibrium of a game where there is no last period. This does not automatically follow from taking the limit of the finite-*T* equilibrium, since that was derived using backward induction. However, an alternative argument based on the stationarity of the game (similar to what was used to go from the odd-period to even-period case) can be used to derive the result.

The striking feature of the alternate-offer bargaining model is that agreement is reached immediately, and there is no delay. Delay would imply inefficiency, since there is a time cost of waiting to complete the transaction. A key assumption in ensuring immediate agreement is that both parties know the valuations of the good – there is no asymmetry of information. As a result, the bargaining process in practice is trivial.

There are many ways of introducing incomplete information that leads to delays in agreement. A simple model is presented here, due to Fudenberg and Tirole (1983). It assumes that the seller does not know the buyer's positive valuation, which can be high (v_H) or low (v_L), but knows that each is equally likely. Assume again that the seller's valuation is zero, and this is known to the buyer. Also, the seller is the only one to make offers – the buyer cannot make counter-offers, only accept or reject the seller's offer. The model has only two periods. The seller therefore chooses offers to make in each period, p_1 and p_2. The buyer may reject the first period offer, either because his valuation is low, or because he thinks he can get a lower offer in the second period. The seller, on the other hand, does not want to always price low initially to ensure an immediate sale, because that may involve losing some potential surplus if the buyer has a high valuation. Thus, the model captures some essential features of bargaining.

For simplicity, we focus on the case where $v_L > v_H / 2$. In a one period case, the seller would therefore offer v_L, which is accepted with certainty. The alternative of offering v_H

would be accepted only half the time, with an expected payoff of $v_H / 2$, which is lower by assumption. In the two-period case, if the second period has been reached, then the seller's posterior probability estimate that the buyer's valuation is v_L is at least half, so the second-period price offer must be $p_2^* = v_L$.

In the first period, the low value buyer will not pay more than v_L, but the high value buyer has a higher reservation price. However, this price is less than v_H, because he can get a lower price by waiting. At the same time, waiting has a cost, captured in the buyer's discount factor, δ_B. Hence the acceptance condition for a high value buyer is

$$v_H - p_1 \geq \delta_B (v_H - v_L) \tag{12}$$

This implies that the maximum first period price is

$$p_1 = \delta_B v_L + (1 - \delta_B) v_H \tag{13}$$

If the seller chooses this price in period 1, her expected profit is

$$\frac{1}{2}[\delta_B v_L + (1 - \delta_B) v_H] + \frac{1}{2}\delta_S v_L \tag{14}$$

On the other hand, the seller could just charge v_L and get that amount as a sure gain. The condition for the higher price being an equilibrium is therefore given by

$$\begin{aligned} &\frac{1}{2}[\delta_B v_L + (1 - \delta_B) v_H] + \frac{1}{2}\delta_S v_L \geq v_L, \quad \text{or} \\ &(1 - \delta_B) v_H \geq (2 - \delta_B - \delta_S) v_L \end{aligned} \tag{15}$$

Given the earlier requirement that $2v_L > v_H$, this cannot hold if the discount factors are equal for the buyer and seller, but is possible if the seller is sufficiently more patient than the buyer. The condition for that to hold is

$$\delta_S \geq [\delta_B(v_H - v_L) + 2v_L - v_H] / v_L \tag{16}$$

Since the seller's discount factor cannot exceed one, a further necessary condition is $1 \geq [\delta_B(v_H - v_L) + 2v_L - v_H] / v_L$, but this is automatically satisfied when the buyer's discount factor does not exceed one.

To summarize, if the seller is sufficiently more patient than the buyer, the equilibrium is

$$(p_1^*, p_2^*) = (\delta_B v_L + (1 - \delta_B) v_H, v_L) \tag{17}$$

In this equilibrium, a high value buyer will accept immediately, while a low value buyer will reject initially, and accept the second period offer.

In the case worked out here, the good is always sold. There is some inefficiency from the delay, as the low value buyer has to wait one period. There is also a distributional effect from the asymmetric information (a positive externality from the low value to the high value buyer), in the sense that the high value buyer gets a lower price than would be possible in the absence of the low value buyer. Suppose instead that $2v_L \leq v_H$. Again, there can be delay. Furthermore, there is some chance that the transaction may not occur at all, which implies greater inefficiency. This case is more complicated, so we do not work it out here. The intuition, however, is that the greater relative value of the high value buyer makes the seller act tougher in her offers, which also leads to the possibility of no agreement.

Middlemen

Models of bargaining typically focus on the details of bargaining protocols, and how those affect equilibrium outcomes. Bargaining models can also be embedded in market contexts, where buyers and sellers search for each other and negotiate while conscious of the presence of other options that might be discovered. In this essay we are not focusing on these search technologies, though they also affect price formation. Instead, we are focusing on the institutions that govern the terms of the transaction. With this in mind, we turn next to looking at the role of middlemen or intermediaries. Middlemen economize on search for buyers and sellers, so their role is tied in to the existence of search alternatives. We begin with the simplest case, however, where buyers and sellers must come to a middleman, before discussing what happens when there is an alternative of searching for direct transactions.

We assume that there are many buyers and many sellers, as follows.[4] There is a single indivisible good which is consumed in one unit quantities only. Buyers have reservation prices, v, which are distributed uniformly on [0,1]. Sellers have reservation values w, also uniformly distributed on the unit interval. A buyer's utility from paying a price p is $U_B(v) = v - p$. The traditional market demand function is $D(p) = 1 - p$, assuming $0 \leq p \leq 1$. A seller's net utility from receiving price p is $U_S(w) = p - w$. The traditional market supply function is $S(p) = p$. Reservation values are private information, known only to each individual. A traditional market equilibrium would be one where supply equals demand, or $p^* = ½$. Recall that traditional theory provides no explanation of how this price is reached.

Suppose that traders can only go to a middleman who quotes fixed bid and ask prices, p_b and p_a. These prices are publicly observed by the whole market (posted prices). If the intermediary attracts q_B buyers and q_S sellers, then its profit is

[4] The model presented here is due to Gehrig (1993). Alternative models of intermediaries that yield different outcomes include Spulber (1996a, 1996b).

$$\pi = \min\{q_B, q_S\}(p_a - p_b) \tag{18}$$

The expression captures the possibility that one side of the market may be rationed – some buyers or sellers will not be able to transact.

Intuitively, the middleman will set the bid and ask prices symmetrically around the competitive equilibrium. This can be proved formally, but we will assume this, so that the number of buyers and sellers is equal. Hence $q_B = q_S$ and $p_a = 1 - p_b$. The middleman's profit is therefore given by

$$\pi = p_b(1 - p_b - p_b) = p_b(1 - 2p_b) \tag{19}$$

By calculating the maximum of this function it is easy to see that $p_b^* = \frac{1}{4}$ and so $p_a^* = \frac{3}{4}$. Substituting into the demand or supply function, we see that the equilibrium quantity is $q^* = \frac{1}{4}$.

In this situation, the monopolist is a traditional monopolist with respect to buyers, but also acts as a monopsonist with respect to sellers. The middleman's profit comes solely from being the only channel for buyers and sellers to transact. In the above model, if there are two competing intermediaries, then the equilibrium essentially reduces to a Bertrand equilibrium, or p^* = ½. The logic of the argument is essentially the same as for the case of two price-setting sellers, except it now applies on both sides of the market.

Another possibility is that, even with one middleman, competition can come from buyers and sellers searching and finding each other. This possibility is now described, and its implications for the middleman's equilibrium worked out. Therefore, suppose that buyers and sellers are matched at random as they search. To keep matters very simple, assume that buyers and sellers always find each other (rather than buyers finding other buyers, or sellers other sellers). When a buyer and seller are matched, one of them at random gets to make a take-it-or-leave it offer. Rejection results on no trade taking place.

Because there is the alternative of going to the middleman, the distributions of buyers and sellers who search may not be the entire distributions. Denote the distribution of searching buyers by $F_B(v)$ and of searching sellers by $F_S(w)$. Buyers and sellers also have bid or ask strategies, denoted by $b(v)$ and $a(w)$. We will assume that these strategies are strictly increasing functions of the reservation values. Hence, we have the following expected utility expressions for buyers and sellers, where the first term is the case where that person bids, and the other is where the counterparty bids.

$$\begin{aligned} U_B(v) &= \frac{1}{2}\int_0^{b(v)}[v-b(v)]dF_S(w) + \frac{1}{2}\int_0^{a^{-1}(v)}[v-a(w)]dF_S(w) \\ U_S(w) &= \frac{1}{2}\int_{a(w)}^{1}[a(w)-w]dF_B(v) + \frac{1}{2}\int_{b^{-1}(w)}^{1}[b(v)-w]dF_B(v) \end{aligned} \tag{20}$$

For the buyer, the expectation is over types of sellers. If the buyer bids, the first part of the integrand is independent of the seller's valuation. If the seller asks, the asking price depends on the seller's valuation. The buyer's bid is accepted if the seller's valuation is below the bid, giving the seller nonnegative utility. The seller's ask is accepted if the buyer's valuation is above the asking price. The logic is similar for the components of the seller's expected utility.

We now argue that the distributions of buyers and sellers who search have a simple form. First, note that buyers who search face the risk of meeting a seller with a high valuation and not being able to complete the deal. On the other hand, buying from the intermediary avoids this risk, but at a cost of paying a possible premium to the middleman. Suppose a buyer with valuation v_0 prefers searching to going to the middleman. Then any buyer with a lower valuation could imitate the v_0 strategy and also do better with search. A detailed proof is involved, but the intuition is that imitating the higher valuation buyer gives some advantage with respect to price in the search market, but not in dealing with the intermediary, who charges the same price regardless. Hence, the set of buyers going to the intermediary will be those with valuations in an interval of the form $(\overline{v},1]$. Furthermore, buyers with low enough valuations will not even bother to search, since they would not even be served in a perfect competitive market. Hence the distribution of buyers who search is uniform on an interval of the form $[\underline{v},\overline{v}]$. Similar arguments imply that the distribution of sellers who search is uniform on an interval of the form $[\underline{w},\overline{w}]$, with those above this interval not searching, and those below choosing to go to the middleman. Roughly, the intermediary is the better option for those who have the most to gain from completing a transaction, because it avoids the possible failure that can occur in search.

Now we derive the buyer's and seller's bid and ask strategies. The part of the buyer's expected utility that she can control is given by

$$\int_0^{b(v)}[v-b(v)]dF_S(w)=\int_{\underline{w}}^{b(v)}\frac{1}{\overline{v}-\underline{w}}[v-b(v)]dw \tag{21}$$

The denominator in the integrand is the range of agents that are active in search. The integral is a simple quadratic function in *b*, and the maximum is given by

$$b^*(v)=(v+\underline{w})/2 \tag{22}$$

A similar chain of reasoning gives the seller's optimal ask function, which is

$$a^*(w)=(w+\overline{v})/2 \tag{23}$$

Using these bid and ask functions, we can write the buyer's expected utility as

$$U_B(v)=\frac{1}{2}\int_{\underline{w}}^{(v+\underline{w})/2}\frac{(v-\underline{w})}{2(\overline{v}-\underline{w})}dw+\frac{1}{2}\int_{\underline{w}}^{2v-\overline{v}}\frac{[v-(w+\overline{v})/2]}{(\overline{v}-\underline{w})}dw \tag{24}$$

At the cutoff level $\overline{v}$, this simplifies to $(\overline{v}-\underline{w})/4$. The borderline seller's utility is identical.

This utility value puts a limit on the price that the middleman can charge. The utility for this marginal buyer, indifferent between searching and going to the intermediary, is $(\overline{v}-p_a)$. Equating the two utility expressions, we have

$$p_a=\frac{3}{4}\overline{v}+\frac{1}{4}\underline{w}=\frac{3}{4}(1-\underline{w})+\frac{1}{4}\underline{w}=\frac{3}{4}-\frac{1}{2}\underline{w} \tag{25}$$

The second equality exploits the symmetry of the distributions of buyers and sellers. We will rewrite the expression as the cutoff seller valuation in terms of the ask price. Furthermore, we will use the symmetry of the ask and bid price around ½, so that $p_a=1-p_b$. Hence we have the following expressions for profit:

$$\pi=\underline{w}(p_a-p_b)=(\frac{3}{2}-2p_a)(2p_a-1) \tag{26}$$

The maximum of this with respect to p_a is easily calculated to be 5/8. Hence, the bid price p_b to buyers is 3/8. The volume of transactions is still ¼, as in the monopoly case, but the presence of the alternative reduces the market power of the intermediary.

The above model is just one example of price-setting by active intermediaries. The key idea is that imperfect alternatives to using a monopoly middleman can limit the intermediary's market power. The result that with two middlemen market power goes away entirely is a special one however, and is sensitive to the specification of the matching and search process. In general, competition among intermediaries will reduce the bid-ask spread, but not eliminate it entirely.

Auctions

Auctions are another variation on institutions used to determine the terms of a transaction. We have examined various institutions for price determination, differentiated by the assumptions about the nature of competition and the information available to participants. The simplest example of an auction is when there is a single unit of a good, a single seller, and two potential buyers. For the auction institution to make sense, or be interesting, it should provide some advantage to the seller. This arises when the seller does not know buyer valuations. In that case, competition between the potential buyers through bidding in the auction benefits the seller.

Even within this simple class of cases, there can be many variations in the details of institutional design, with respect to the rules of bidding, constraints on participation, and assumptions about the relationship of buyer valuations. We illustrate the simplest case again, where buyer valuations are independent (and hence uncorrelated), there is no entry condition or participation fee, and bids are made through a simultaneous sealed-bid method. This latter assumption contrasts with the more popular method of repeated open bidding, as featured in art auctions, or in its electronic form on eBay.

Suppose that buyer valuations, denoted by v_1 and v_2, are drawn independently from the same distribution.[5] To keep matters simple, let this distribution be uniform on [0,1]. Both buyers and the seller know this fact (and, more strongly, it is common knowledge), but only the buyer knows his own valuation. The seller's valuation is zero, so there are potential gains to completing the auction successfully.

Let b_i denote the bid made by buyer *i*. The first auction institution we consider is where the highest bidder wins, and pays that winning bid to complete the transaction. The winning bid is therefore the transaction price for the good being auctioned.

It is convenient to assume that buyers use strategies of the form

[5] The following treatment of auctions is standard, and very similar to that presented in Mas-Colell et al. (1995).

$$b_i = k_i v_i \quad \text{for} \quad k_i \in [0,1] \tag{27}$$

This reduces the choice of strategies from the more complicated problem of picking a function of the valuation to picking the parameter k_i. Consider buyer 1's expected utility maximization problem, when buyer 2 chooses a strategy of the above form:

$$Max_{\ b_1 \geq 0} \quad (v_1 - b_1)Prob(b_2(v_2) \leq b_1) \tag{28}$$

The first term is the gain in utility from winning the auction. The second term is the probability of winning, i.e., the probability of outbidding buyer 2. Since bids are submitted simultaneously, buyer 1 has to estimate what buyer 2's bid will be – he cannot observe the rival's bid. Using the strategy functional form in (27), the condition for winning is $v_2 \leq (b_1 / k_2)$. At the same time, buyer 1's bid should never exceed k_2, since this is the maximum that buyer 2 will ever bid. Buyer 1's problem is therefore

$$Max_{\ k_2 \geq b_1 \geq 0} \quad (v_1 - b_1)(b_1 / k_2) \tag{29}$$

The solution is

$$b_1(v_1) = \min\{\tfrac{1}{2} v_1, k_2\} \tag{30}$$

In other words, $k_1 = \frac{1}{2}$ as long as k_2 is not too small.

By similar reasoning, using the symmetry of the buyers,

$$b_2(v_2) = \min\{\tfrac{1}{2} v_2, k_1\} \tag{31}$$

It is easy to see now that $k_1 = k_2 = \frac{1}{2}$ yields a Bayesian Nash equilibrium of the bidding game between the two buyers. In other words, the equilibrium strategies are defined by

$$b_i(v_i) = \tfrac{1}{2} v_i \tag{32}$$

Each buyer only bids one-half of his valuation in this equilibrium. The intuition is as follows. Bidding one's valuation increases the chance of winning the auction, but reduces the surplus from acquiring the good. Hence it makes sense to shade one's bid. In the case of the uniform distribution of valuations, this shading is precisely half of the true valuation. The buyer with the higher valuation will always win the auction. The seller's revenue from this equilibrium is $\max\{\tfrac{1}{2}v_1, \tfrac{1}{2}v_2\}$. Expected revenue therefore is given by

$$E(R_I) = \int_0^1 \left[\int_{v_2}^1 \tfrac{1}{2} v_1 dv_1 + \int_0^{v_2} \tfrac{1}{2} v_2 dv_1 \right] dv_2 \tag{33}$$

The double integration takes account of the seller's uncertainty about both buyers' valuations. The independence of valuations is reflected in the independent distribution functions. The first interior integral captures possibilities where buyer 1 has the higher valuation and higher (winning) bid, the second where buyer 2 wins. This expression can be evaluated, and turns out to be 1/3. If the seller knew buyer valuations, she could extract the entire surplus from the winning bidder through competition, and the expected revenue would be ½ (the mean of the distribution of either valuation). Hence the seller's lack of perfect knowledge of buyer valuations limits her market power to some extent.

Now consider a variation of the above auction. Suppose it is still a sealed-bid auction, and the highest bid wins, but now the winning bidder only pays an amount equal to the second-highest bid. This kind of auction is called a *second-price auction* (the previous one being a *first-price auction*). This method of determining the transaction price does away with any incentive that buyers might have to shade their bids. Therefore, the equilibrium bidding strategies now have the form $b_i(v_i) = v_i$. This is the Bayesian Nash equilibrium of the bidding game between the buyers. In fact, the strategies in this case are weakly dominant (stronger than a Bayesian Nash equilibrium). The buyer with the higher valuation always wins, and this probability is ½, by symmetry.

The seller's expected revenue in this case is given by the expression:

$$E(R_{II}) = \int_0^1 \left[\int_{v_2}^1 v_2 dv_1 + \int_0^{v_2} v_1 dv_1 \right] dv_2 \tag{34}$$

It is important to see that, in comparing this with previous revenue expression, the coefficients of ½ have disappeared, but the valuations have also switched their places in the integrals. The bids are higher, but the seller gets lower payments since she gets only the second-highest bid. If we evaluate the integral, it turns out to be 1/3 again. This is not a coincidence. The equality is a special case or illustration of the *revenue equivalence theorem*. The key assumptions of the theorem are risk neutral buyers with independent valuations drawn from the same distribution. Then if two auctions give the lowest valuation buyer the same expected utility, and probabilities of winning are the same for the same vector of realized valuations, then the seller's expected revenue must be equal across the auctions.

As noted earlier, there are many variations of basic auction institutions, and assumptions and outcomes can differ. The general lesson to take away is that auctions can use competition among potential buyers to extract more revenue from the eventual buyer than might be possible in a bilateral bargaining situation. It is also possible that sellers compete as well. Two potential sellers might simultaneously run auctions for two potential buyers. This would tend to increase the leverage of buyers. On the other hand, if the sellers' valuations are positive and unknown, the buyers might risk not getting the item at all – this could increase the market power of the sellers. In sum, many complications in market institutions can be envisaged, with the competitive model of price-taking behavior providing a benchmark for all of them.

References

Fudenberg, Drew, and Jean Tirole, 1983, Sequential Bargaining with Incomplete Information, *Review of Economic Studies*, Vol. 50, No. 2, pp. 221-247

Gehrig, Thomas, 1993, Intermediation in Search Markets, *Journal of Economics & Management Strategy*, Vol. *2,* No. 1, pp. 97-120

Kreps, David, 1990, *A Course in Microeconomic Theory*, Princeton, NJ: Princeton University.

Manning, R. and P. B. Morgan,1982, Search and Consumer Theory, *The Review of Economic Studies*, Vol. 49, No. 2, pp. 203-216.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, 1995, *Microeconomic Theory,* New York: Oxford University Press.

Singh, Nirvikar, 2010a, "Uncertainty and Information: An Expository Essay," University of California, Santa Cruz, Department of Economics.

Riley, John, and Richard Zeckhauser, 1983, Optimal Selling Strategies: When to Haggle, When to Hold Firm, *Quarterly Journal of Economics*, Vol. 98, No. 2, pp. 267-289

Rothschild, Michael, and Joseph Stiglitz, 1976, Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information, *Quarterly Journal of Economics,* Vol. 90, No. 4, pp. 629-649

Rothschild, Michael, 1973, Models of Market Organization with Imperfect Information: A Survey, *Journal of Political Economy*, Vol. 81, No. 6, pp. 1283-1308

Rubinstein, Ariel, 1982, Perfect Equilibrium in a Bargaining Model, *Econometrica*, Vol. 50, pp. 97-109

Shaked, Avner, and John Sutton, 1984, Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model, *Econometrica*, Vol. 52, pp. 1351-1364

Singh, Nirvikar, 2010b, “Transaction, Search and Switching Costs: An Expository Essay,” University of California, Santa Cruz, Department of Economics.

Spulber, Daniel F., 1996a, Market Microstructure and Intermediation, *Journal of Economic Perspectives*, Vol. 10, No. 3, pp. 135-152

Spulber, Daniel F., 1996b, Market Making by Price-Setting Firms, *Review of Economic Studies*, Vol. 63, No. 4, pp. 559-580

Varian, Hal R., 1992, *Microeconomic Analysis*, 3rd ed., New York: Norton.