



Edinburgh School of Economics
Discussion Paper Series
Number 86

The provision of quality in a bilateral search market

Clara Ponsati (IAE (CSIC), Barcelona)
Jozsef Sakovics (University of Edinburgh)

Date
August 2002

Published by

School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361

<http://www.ed.ac.uk/schools-departments/economics>



THE UNIVERSITY *of* EDINBURGH

The provision of quality in a bilateral search market

Clara Ponsatí – IAE (CSIC), Barcelona

and

József Sákovics – University of Edinburgh

August 2002

Abstract

We accomplish two goals. First, we provide a non-cooperative foundation for the use of the Nash bargaining solution in search markets. This finding should help to close the rift between the search and the matching-and-bargaining literature. Second, we establish that the diversity of quality offered (at an increasing price-quality ratio) in a decentralized market is an equilibrium phenomenon – even in the limit as search frictions disappear.

JEL Classification: C78, L11.

Key words: quality dispersion, Nash Program, bilateral search

Correspondence to: József Sákovics, Department of Economics, University of Edinburgh, 50 George Square, EH8 9JY, Edinburgh. E-mail: Jozsef.Sakovics@ed.ac.uk

1 Introduction

It is a ubiquitous observation in our society that not only the best professionals can make a living. There are plenty of mediocre lawyers, dentists and economists(!) who are clearly outclassed by their peers but still find employment. Of course, there are a number of straightforward explanations of why this should be the case. First, it is likely that the demand for the services of these professionals is heterogenous: some people are willing (and able) to pay more for better service, so if price discrimination is not possible we must have quality diversity. A second argument relates to the scarcity of high quality in the market. If a lower quality is instantaneously available, while for the high quality one must queue, again there is room for lemons to enter the market. A third argument is based on asymmetric information: if at the time of contracting the quality is not observable by the customer, it actually may be the high quality services that are driven out of the market.

While all the above arguments are convincing in their own right, we contend that “quality dispersion” is a natural phenomenon in a decentralized market even if the consumers are identical, the market is frictionless, and the consumers can tell the quality of service before they purchase it.

We show our result in the context of a search/matching market in its steady state. The literature on the analysis of such markets was begun by Diamond (1971), and has evolved a great deal since. There have developed two main strands: search theory emphasizes the endogenous nature of the effort put into search activity, (over)simplifying the exact nature of negotiation between a seller and a buyer – generally, by using the

Nash bargaining solution –; matching-and-bargaining theory does the opposite, it (over)simplifies search behavior – by assuming that it is governed by an exogenous matching process – while it concerns itself more with the details of bargaining. We start out in the spirit of the latter approach, but – rather reassuringly – end up with a model that supports the reduced form approach of search theory.

The original argument against using the Nash bargaining solution in a non-cooperative model of market interaction (mainly due to Rubinstein and Wolinsky (1985), Gale (1986) and Wolinsky (1987)) was simply based on the desire for a fully non-cooperative treatment. A more devastating blow was the discovery of the “Outside Option Principle” (by Ken Binmore, Avner Shaked and John Sutton¹), which has established that the outside options should only have a direct effect on the outcome of bargaining if they exceed the equilibrium payoff in the absence of an option. It was Bester (1988) who incorporated the “full-fledged” non-cooperative bargaining model into a matching market. Thus, in Bester’s world when a buyer and a seller are matched, they engage in alternating-offer bargaining, where the buyer – but only the buyer – has the option of breaking up the negotiations, following his rejection of any offer by the seller.

While in the original set-up of Diamond (1971) – where the sellers post prices and the buyers visit the stores – this seems to be the appropriate model, in the context of pairwise matching and bargaining of, say, professionals and their clients, it is not. It

¹See Binmore (1985), Binmore, Shaked and Sutton (1989), Shaked and Sutton (1984) and Sutton (1986).

is not, because it lacks symmetry. There is no reason whatsoever to rule out that a lawyer refuses to work with a client who does not accept to pay the fee she demands. Consequently, we need to incorporate into the bargaining model the option that both players can leave the negotiating table. There are two obvious ways this can be done. The first one is to allow only the responder to opt out in every period. While, by the alternating structure of the bargaining procedure, this would ensure that both players enjoy the option of quitting, one of them would be forced to suffer some (delay) cost before being able to do so. This does not seem quite reasonable. Thus, we opt for the alternative specification, where both players can leave following the rejection of an offer.²

As we argue it in detail in Section 3, the above modeling choice leads us to a non-cooperative solution, which coincides with the (generalized) Nash bargaining solution.³ Thus, we arrive at the paradoxical(?) conclusion that taking the strategic approach “all the way”, brings us back to the cooperative solution. (A fitting way to culminate the Nash Program – the quest for a non-cooperative foundation of the Nash bargaining solution.)

In addition to the conceptual breakthrough discussed above, we also arrive at

²Actually, the issue is not this simple. The crucial question is whether is it feasible that the proposer can leave the negotiation following the rejection of his offer without listening to the responder’s counter-offer (c.f. Shaked 1994). If (and only if) the answer is yes, our set-up is the appropriate model.

³It has been shown (see Binmore, Rubinstein and Wolinsky (1986)) that the limit of the Rubinstein solution as the players become increasingly patient is the Nash solution. Our equivalence result does not depend on taking any limit.

practical predictions about the market equilibrium, in the presence of a (potentially) disperse distribution of quality.⁴ Our main findings are two-fold. First, – and contrary to Bester (1988) – we establish that the equilibrium price distribution is such that the mark-up of the seller is always increasing with the quality of the service sold. This is in clear accordance with the stylized fact that in general the profit margin is higher for higher valued goods/services. Second, we show that while it is true that as search frictions diminish the average provision of quality increases – just as predicted by Bester (1988) –, in the limit as these frictions disappear the equilibrium does not converge to the degenerate case, where only the highest quality sellers can stay in the market – as predicted by Bester (1988) –, rather we are left with a very significant proportion of mediocre sellers in business. That is, we support our contention mentioned at the beginning of this Introduction that quality dispersion is an “innate” characteristic of decentralized markets.

The rest of this article continues as follows. In Section 2, we lay out the details of our model. In Section 3, we develop our bargaining solution. Section 4 contains the analysis of the market equilibrium and its dependence on the level of search frictions. Section 5 concludes.

⁴Surprisingly – to the best of our knowledge – no search model using the Nash solution has analyzed this question. An exception is Bester (1993), but he only allows for two levels of quality (while makes the firms’ choice of quality endogenous).

2 The model

We consider a market for a single commodity of heterogeneous quality, composed of a set of producers (sellers) and a set of consumers (buyers).⁵ Both sets of agents are assumed to be a continuum and their measure is normalized to one. Correspondingly, each seller can be uniquely indexed by a type, $\theta \in [0, 1]$. Seller θ can produce the good of quality $q(\theta)$, where $q(\cdot) : [0, 1] \mapsto [0, 1]$ is – without loss of generality – assumed to be non-decreasing. For simplicity, we also assume that $q(\cdot)$ is continuous and we normalize⁶ $q(1)$ at 1.

Each producer can sell a single unit in each period. The cost of production is independent of the quality⁷ and, for simplicity, it is normalized to zero. The buyers are all identical. Each of them wishes to purchase a single unit of the good. Their valuation of a good of quality $q(\theta)$ is equal to the quality measure.

The market opens at $t = 0$ and it operates over time. The agents maximize their expected utility, which they discount by the common discount factor, $\delta \in (0, 1)$, per unit of time. Players only receive utility if they consummate a transaction. The utility of a consumer who purchases the product of seller θ for a price of $p(\theta)$ at time t is given by $[q(\theta) - p(\theta)] \delta^t$, while seller θ 's utility gain from the same transaction is

⁵While our principal application is professional services, we will stick to the standard terminology of producers and consumers (of goods).

⁶Since we are not concerned with costs (they are sunk), if the highest quality produced were not 1, we could just shift up the quality distribution, without any real consequence.

⁷Since our main point is that mediocre quality is produced in equilibrium, by not giving mediocre producers a cost advantage, we actually strengthen our result.

$p(\theta) \delta^t$.

The consumers only know the distribution of quality in the market. The game starts by the buyers' search for producers. Each buyer chooses a seller at random and learns the quality of her product upon entering the store. Search is costly. The expected delay cost of finding an empty store is represented by a discount of δ_b . As we will see, the value of this friction is endogenously determined, since it depends on the ratio of active sellers to buyers. Once a consumer finds himself in a store he starts bargaining over the price with the seller.

The bargaining procedure is an enriched version of the standard model of alternating offers. The added feature is that after a rejection, both players are allowed to unilaterally terminate negotiations and return to the market. More specifically, bargaining between matched sellers and buyers proceeds as follows. First, one of the parties is randomly selected to make the first proposal. The probability that the buyer (the seller) is selected is λ ($1 - \lambda$). If the responder accepts the proposed price the transaction is consummated, the agents collect their payoffs, the seller returns⁸ to the market and the buyer is replaced by a new, but unmatched, buyer. On the other hand, if the responder rejects then either of the two bargainers has the option to break up negotiations. In case of a break-up, it "costs" δ_b and δ_s for the buyer and the seller, respectively, to find a new match. If neither player opts out, the responder makes a counter-proposal after a delay of one time unit. The seller and the buyer al-

⁸We assume that the sellers are myopic, in the sense that they do not consider their future flow of income when negotiating with a buyer. Equivalently, they could be replaced by a new seller of the same quality.

ternate in making proposals until either a proposal is accepted or one of them breaks up negotiations. If they continue to bargain forever, they both earn 0.

Note that our market is stationary, since both the measure and the type distribution of the agents are unaltered over time. Thus, in equilibrium, the expected gain of a consumer upon entering the market, v , is constant. Since this value is also the continuation value of a buyer upon leaving a store empty-handed, when it is larger than the minimum quality produced, there are producers who are unable to sell in equilibrium. We denote the marginal producer by θ^* and we assume that in equilibrium all producers with lower quality are absent from the market.⁹ Thus, an equilibrium of this market can be described by θ^* and the expected price¹⁰ at each producing seller's store: $p(\theta)$ for $\theta \geq \theta^*$.

Let us return to the search frictions. To capture the dependence of these on θ^* , we assume that for each buyer the probability of a match with a seller is $1 - \theta^*$ – the number of sellers per buyer – in each period¹¹ of (memoryless) matching. Let us denote the (common) discount factor corresponding to one matching period by δ^Δ . The effective discount factors of a seller and a buyer are then determined by the

⁹Alternatively, we could assume that the non-trading sellers are still hanging around and thus it is possible to be matched to them. The resulting search frictions would be identical (as long as the buyers remain the long side of the market).

¹⁰As we will see, in equilibrium every match ends with immediate agreement.

¹¹Think of a matching period as the time it takes for a seller (on the short side of the market) to find a buyer.

following equations¹²

$$\delta_s = \delta^\Delta \text{ and } \delta_b = \delta^\Delta (1 - \theta^* + \theta^* \delta_b). \quad (1)$$

Resolving the equations, we obtain

$$\delta_s = \delta^\Delta \text{ and } \delta_b = \frac{\delta^\Delta (1 - \theta^*)}{1 - \delta^\Delta \theta^*}. \quad (2)$$

3 The bargaining solution

We start our analysis by considering the negotiation between a buyer-seller pair in isolation. To this effect, in this section, we assume that the outside options are exogenously given. Let x_s and x_b denote the seller's and the buyer's outside option respectively. First, we characterize the set of subgame-perfect equilibria of the subgames where one of the parties has already been chosen to be the first proposer. Let us denote the equilibrium price of the subgame where the buyer (the seller) is the first proposer by $p_b^*(\theta)$ ($p_s^*(\theta)$).

Proposition 1 i) *If $x_s + x_b > q(\theta)$, then the unique equilibrium outcome is an instant break-up of negotiations in both subgames.*

ii) *Otherwise, immediate agreement at the prices $p_b^*(\theta) = x_s$ and $p_s^*(\theta) = q(\theta) - x_b$ can be supported by subgame-perfect equilibria independently of the values of the outside options and of the players' time preferences.*

¹²For readers more familiar with the continuous time formulation (Bellman equations), we provide the equivalent treatment in the Appendix.

iii) Finally, if $x_i \leq \delta(\delta q(\theta) - x_{-i})$ for $i = s, b$, then there exist a continuum of equilibrium outcomes. First, immediate agreements at $p_b^*(\theta) \in [x_s, \delta q(\theta) - x_b]$ and $p_s^*(\theta) \in [q(\theta)(1 - \delta) + x_s, q(\theta) - x_b]$, respectively. In addition, there exist delayed agreements at a subset of these same prices.

Proof: It is a straightforward adaptation of Lemma 1 and the Theorem of Ponsatí and Sákovics (1998) to the present model. •

When the aggregate value of the outside options exceeds the gains from trade, it is evident that at least one of the traders will prefer not to trade. There is one equilibrium which always exists, whenever it is socially optimal to trade. This equilibrium, henceforth the *equilibrium in ultimatum strategies*, is sustained by a credible threat of the proposer to leave the bargaining table in case his offer is rejected. The credibility of this threat comes from the similar threat that the current responder makes in the following period when she will be the proposer. To see this, imagine for a second that the outside options are valueless. In this case, the equilibrium in ultimatum strategies would give the entire surplus to the first proposer. Note that the threats are sustaining each other: given that the proposer expects no gains tomorrow, his threat of leaving today is credible. When the outside options are of positive value, this argument is modified to the extent that the responder will have to be given at least her outside option to be willing to trade.¹³

¹³Note that a positive outside option for the first proposer does not affect the argument, since his payoff upon quitting rises together – in fact, even faster, because of discounting – with his expected

As in any version of strategic bargaining models, in equilibrium the responder has to be indifferent between accepting or not, since otherwise the proposer could increase his share. In the equilibrium in ultimatum strategies, this indifference is between accepting and opting out. This leaves open another possibility to construct equilibria: one where the responder is indifferent between accepting and continuing to the next period (and prefers both to opting out, of course). Since both of these values are determined endogenously, it is not surprising that there is a continuum of ways that this indifference can be achieved. Note also, that for this type of equilibrium to exist, we also need that the proposer be willing to continue in case his offer is rejected. Finally, as it is standard in bargaining games which have multiple efficient equilibria, a threat to switch to an extreme equilibrium can support delayed agreements at some intermediate prices.

The equilibrium in ultimatum strategies is salient, because it has a number of attractive features.

i) As it is apparent from the statement of the result, this is the only equilibrium which *always exists* when there are gains from trade.

ii) By the same token, when the outside options are sufficiently large compared to the discount factor (if $\max(\frac{x_s}{\delta} + \frac{x_b}{\delta^2}, \frac{x_s}{\delta^2} + \frac{x_b}{\delta}) > q(\theta)$) this equilibrium is *unique*.

iii) Also, this is the only equilibrium in which strategies are *independent of the players' time preference*.¹⁴ As a consequence, our equilibrium is valid even

payoff if he continues to bargain.

¹⁴This result generalizes to the case when the players have different discount factors.

when the players are uncertain about each others' discount factor.

iv) Finally, it is easy to see that this equilibrium outcome is the *unique equilibrium outcome for any finite horizon truncation* of the bargaining game. As a consequence, our equilibrium is valid (and unique) even when there is a deadline, before which agreement must be reached.

In view of the remarkable robustness of the equilibrium in ultimatum strategies, we propose that it be considered as the “predicted behavior” in the bargaining problem. That is, we assume that

Assumption A. 1 *The negotiation between a buyer and seller θ results in immediate agreement, with the expected price given by*

$$p(\theta) = \lambda x_s + (1 - \lambda)(q(\theta) - x_b), \quad (3)$$

whenever there exist gains from trade ($x_s + x_b \leq q(\theta)$).

It is remarkable that this bargaining solution coincides with the asymmetric Nash solution (see, Harsányi and Selten (1972)), where the outside options are interpreted as the disagreement outcome, while λ and $1 - \lambda$ are the bargaining weights of the seller and the buyer, respectively.

Some readers may doubt the attractiveness of this feature in a model with impatient players. Note, however, that in the full model of the decentralized market the effects of the players' impatience while searching does filter in through the (endogenous) outside options.

4 The market equilibrium

In order to characterize our market equilibrium we need to do two things (simultaneously). Adapt the bargaining solution defined in the preceding section to the endogenous outside options of our search market and determine the marginal quality produced.

Except for the cost of waiting, a seller's outside option does not vary from her current expected value, which is given by the expected price:

$$x_s(\theta) = \delta_s p(\theta) \tag{4}$$

From (3) and (4) we obtain

$$x_s(\theta) = \frac{(1 - \lambda)\delta_s}{1 - \lambda\delta_s} \cdot (q(\theta) - x_b). \tag{5}$$

Since the coefficient is less than unity, (5) implies that the necessary and sufficient condition for the existence of gains from trade – and, therefore, trade – is $q(\theta) \geq x_b$. Of course, the buyers' outside option depends on the distribution of quality offered. Consequently, the lowest quality producer, θ^* , is defined by $q(\theta^*) = x_b(\theta^*)$ – or by zero when $q(0) \geq x_b(0)$.¹⁵

In order to determine the marginal quality, we need to calculate the buyers' outside option. This is equal to the discounted expected profits from a future match. Note

¹⁵This definition applies for any marginal producer satisfying the (in)equality, since – once θ^* is fixed – the outside option is constant, while the quality produced is increasing in θ . Consequently, potentially there could be multiple solutions.

that, since the matching technology is memoryless, this expected value is independent of the quality of the good they are currently bargaining for. Instead, it is a function of the distribution of quality produced. Namely,

$$x_b = \frac{\delta_b}{1 - \theta^*} \cdot \int_{\theta^*}^1 [q(\theta) - p(\theta)] d\theta \quad (6)$$

Using (4), (5) and (6), we obtain

$$x_b = \frac{\lambda\delta_b(1 - \delta_s)}{1 - \lambda\delta_s - \delta_b(1 - \lambda)} \cdot AQ(\theta^*) = \frac{\lambda\delta^\Delta}{1 - \lambda\delta^\Delta\theta^*} \int_{\theta^*}^1 q(\theta) d\theta, \quad (7)$$

where $AQ(\theta^*) = \int_{\theta^*}^1 \frac{q(\theta)}{1 - \theta^*} d\theta$ denotes the average quality produced in equilibrium.

Theorem 2 *If and only if*

$$q(0) \geq \lambda\delta^\Delta \int_0^1 q(\theta) d\theta, \quad (8)$$

all the sellers will produce. In this case the equilibrium prices are given by

$$p(\theta) = \frac{1 - \lambda}{1 - \lambda\delta^\Delta} \cdot \left(q(\theta) - \lambda\delta^\Delta \int_0^1 q(\theta) d\theta \right). \quad (9)$$

When (8) is not satisfied, the marginal seller is uniquely determined by the solution to

$$\frac{\int_\theta^1 q(x) dx}{q(\theta)} = \frac{1}{\lambda\delta^\Delta} - \theta, \quad (10)$$

while the corresponding prices are given by

$$p(\theta) = \frac{1 - \lambda}{1 - \lambda\delta^\Delta} \cdot [q(\theta) - q(\theta^*)].$$

Proof: By the above argument, we will have $\theta^* = 0$, when $q(0) \geq x_b$, or (8). In this case, all the sellers will produce and the equilibrium price function follows from (4), (5) and (7) evaluated at $\theta^* = 0$.

If (8) is not satisfied, then $x_b = q(\theta^*)$. Substituting into equation (7) and simplifying we obtain the equilibrium condition

$$H(\theta^*) = \frac{1}{\lambda\delta\Delta} - \theta^*, \quad (11)$$

where $H(\theta) = \frac{\int_{\theta}^1 q(x) dx}{q(\theta)}$. Observe that, in the interval $[0, 1]$, $H(\theta)$ is continuous and it is monotonically decreasing, from $\frac{\int_0^1 q(x) dx}{q(0)} \geq 1$ to 0. In addition, $H'(1) = -1$. On the other hand, the RHS of (11) is larger than $H(0)$ at $\theta = 0$ – by the fact that (8) is not satisfied –, but it is positive (and therefore larger than $H(1)$) at $\theta = 1$. Consequently, there always exists an interior solution. Thus, the existence of a market equilibrium is guaranteed.

To see uniqueness (for both cases), just note that the slope of $H(\theta)$ is $-1 - \frac{q'}{q}H(\theta)$, which is strictly less than -1 , for $\theta < 1$. •

insert *Figure1*

Corollary 3 *In equilibrium in every match there are always strictly positive gains to trade: $q(\theta^*) > 0$.*

Proof: When $q(0) > 0$, the result is obvious. Otherwise, $H(\theta) = \infty$, for all θ such that $q(\theta) = 0$, so $q(\theta^*)$ must be positive. •

The fact that the minimum quality provided is strictly above the buyers' threshold

level has an important effect on the equilibrium “price-quality ratio”. Note that

$$\frac{p(\theta)}{q(\theta)} = \frac{1 - \lambda}{1 - \lambda\delta^\Delta} \cdot \left[1 - \frac{\max \left\{ q(\theta^*), \lambda\delta^\Delta \int_0^1 q(\theta) d\theta \right\}}{q(\theta)} \right], \quad (12)$$

which is strictly increasing in $q(\theta)$. That is, the proportion of the gains from trade appropriated by the sellers is increasing with the quality of their good. In other words, the mark-up is increasing with the quality, a commonly observed fact of everyday life.

We can now determine how the marginal level of quality produced varies as the common search frictions change. Note that the LHS of (10) is independent of δ^Δ , while the RHS is (uniformly) decreasing in it. Consequently, θ^* is increasing in δ^Δ . That is, we confirm that even in the case when the sellers (as well as the buyers) can terminate negotiations, increasing the common search frictions decreases the average quality sold in the market. On the other hand, in the limit as search frictions disappear,¹⁶ we have

$$\frac{\int_{\theta^*}^1 q(x) dx}{q(\theta^*)} = \frac{1}{\lambda} - \theta^*, \quad (13)$$

which yields $\theta^* < 1$, whenever $\lambda < 1$. Thus, we have shown that

Corollary 4 *In the (asymptotically) frictionless market, a significant range of qualities is provided. The level of quality dispersion is increasing in the sellers’ bargaining power.*

¹⁶In fact, given that the application we have in mind is in the service sector, it is plausible that the provision of the (private) service takes time. That is, even as search frictions disappear, there is still inherent delay in the market. It would be easy to model this explicitly but there is no point in doing it, since it would only bias the marginal quality (close to the limit) further downward.

This is the main result of this analysis. Note that the level of quality dispersion is significant. For example, under the assumption that quality is uniformly distributed between zero and one, it is easy to show that

$$\theta^* = \frac{1 - \sqrt{1 - \lambda^2}}{\lambda}.$$

As we can see in Figure 2 below - plotting θ^* as a function of λ - unless the buyers have almost all the bargaining power, the quality dispersion in the market will be very significant.

insert *Figure 2*

5 Conclusions

In this paper, we have provided non-cooperative foundations for the (generalized) Nash Bargaining Solution, which are specifically applicable to market models of decentralized negotiation. The strategic form that we propose for the process of bilateral bargaining witnesses the large amount of flexibility the parties have while negotiating. In particular, we assume that each party can voluntarily and credibly leave the negotiation (for good) following the rejection of any offer.

When we incorporate the Nash Bargaining Solution into a decentralized market with asymmetric information about quality, we obtain that – in equilibrium – the share of the gains from trade that accrues to the sellers is increasing with the quality of the sellers' product. This result does not obtain using the alternative strategic

bargaining models, whose solutions are not directly related to the outside options.

Additionally, we also overturn the result of Bester (1988), where he obtains that as search frictions disappear only the highest quality seller will produce. The result that in our model there is quality dispersion – even in the limit – is a consequence of carefully modelling the sellers’ option of refusing to sell to a given customer. In particular, we have taken account of the fact that the search cost facing a buyer is closely related to the number of sellers in the market. That is, as we decrease the length of a matching period, the ratio of the buyers’ to sellers’ search costs is increasing. Since the marginal quality is determined by the buyers’ continuation value this provides a force, which countervails the buyers’ increased willingness to wait for a better match, which balances out at an intermediate quality level.

The question that begs to be asked is how efficient the equilibrium provision of quality is. There is not a straightforward answer to this question, for two reasons. First, it is unclear what should the benchmark be. A naive view could say that we should maximize the amount of surplus generated per period, so the optimal solution is for all sellers to produce. However, this solution does not take into account that a buyer may prefer to wait rather than trade with a low quality seller. Second, since by varying the bargaining weights the value of the marginal quality spans the entire range, no matter what the benchmark is, by adjusting the relative bargaining power we can always match the outcome. Thus, unless we want to calibrate the model and thus hypothesize a true distribution of bargaining power, a welfare analysis is of limited interest.

Appendix

Here we provide the derivation of the buyers' Bellman equation, for the equivalent specification in continuous time. Denote by V and H a buyer's continuation value when unmatched and matched, respectively. Denote the common interest rate by r and the "arrival rate" of sellers by s . For this treatment is advantageous to assume that this arrival rate is of any seller, and thus matches with seller index below θ^* are unsuccessful. We then have (c.f. equation (1))

$$V = e^{-r\Delta} (s\Delta(1 - \theta^*)H + (1 - s\Delta(1 - \theta^*))V).$$

Solving for V , and approximating $e^{-r\Delta}$ by $1 - r\Delta$ (and ignoring terms including Δ^2), we obtain (c.f. equation (2))

$$V = H \frac{s(1 - \theta^*)}{r + s(1 - \theta^*)}.$$

If we take into account that $V = q(\theta^*)$ –as argued in the paper– we can rewrite the above as

$$rV = s(1 - \theta^*)(H - q(\theta^*)),$$

which may look more familiar.

References

- [1] Bester, H. (1988), "Bargaining, Search Costs and Equilibrium Price Distributions," *Review of Economic Studies*, **55**, 201-214.

- [2] Bester, H. (1993), "Bargaining versus Price Competition in Markets with Quality Uncertainty," *American Economic Review*, **83** (1), 278-288.
- [3] Binmore, K. (1985), "Bargaining and Coalitions," in A. Roth (ed.) *Game Theoretic Models of Bargaining*, Cambridge University Press.
- [4] Binmore, K., Rubinstein, A. and A. Wolinsky (1986), "The Nash Bargaining Solution in Economic Modelling," *RAND Journal of Economics*, **17**, 176-188.
- [5] Binmore, K., Shaked, A. and J. Sutton (1989), "An Outside Option Experiment," *Quarterly Journal of Economics*, **104**, 753-770.
- [6] Diamond, P. (1971), "A Model of Price Adjustment," *Journal of Economic Theory*, **3**, 156-168.
- [7] Gale, D. (1986), "Bargaining and Competition Part I: Characterization," *Econometrica*, **54**, 785-806.
- [8] Harsányi, J. and R. Selten (1972), "A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information," *Management Science*, **18**, 80-106.
- [9] Ponsatí, C. and J. Sákovics (1998), "Rubinstein Bargaining with Two-Sided Outside Options," *Economic Theory*, **11**, 667-672.
- [10] Rubinstein, A. and A. Wolinsky (1985), "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, **53**, 1133-1150.

- [11] Shaked, A. and J. Sutton (1984), “Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model,” *Econometrica*, **52**, 1351-1364.
- [12] Sutton, J. (1986), “Non-Cooperative Bargaining Theory: An Introduction,” *Review of Economic Studies*, **53**, 709-724.
- [13] Wolinsky, A. (1987), “Matching, Search and Bargaining,” *Journal of Economic Theory*, **42**, 311-333.




Figure 1: A typical layout of the curves determining the marginal producer




Figure 2