The Power of a Strategic Buyer with Perfect and Imperfect Monitoring

V. Bhaskar
Dept. of Economics
University of Essex
Wivenhoe Park
Colchester CO4 3SQ, UK.
Email: vbhas@essex.ac.uk

June 2000

Abstract

We analyze the implications of the buyer's strategic power in the model of dynamic price competition. Two sellers with a limited inventory sell to a single buyer, who has unit for in each of several time periods. The market power of the sellers is offset by the strategic power of the buyer. By not consuming in any period, the buyer can destroy a unit of demand, thereby intensifying future price competition. If transactions are perfectly observed, we find that that a strategic buyer can do significantly better than non-strategic buyers. This excercise of strategic power may give rise to inefficiencies. However, if an agent only perfectly observes those transactions in which he is directly involved, and imperfectly observes other transactions, the strategic power of the buyer is reduced. In some cases, buyer power may be completely eliminated.

1 Introduction

We analyze the implications of the buyer's strategic power in the model of dynamic price competition. Two sellers, A and B, each have a finite endowment of inventory, ω_A and ω_B. They sell to a single buyer, C, who demands a single unit of the product in each of n periods, and whose valuation in every period of this single unit is normalized to one. The sellers set prices, and hence have market power. Our model is relevant for agricultural products such as coffee — producing countries such as Colombia can exert market power via there policies, but buyers such as Nestlé are also large.

Our work builds upon Dudey (1992), who analyzed this model of dynamic Bertrand-Edgeworth competition when the buyers in different periods are different agents (or equivalently, when the single buyer is myopic). His main conclusion was that if the sellers differ, so that $\omega_A > \omega_B$, and if the smaller of the two sellers could not meet the entire demand ($\omega_B < n$), the sellers would exercise effective monopoly power, and the buyer(s) would obtain no share of the surplus. Furthermore, the smaller of the two sellers, B, has more strategic power — he would be able to sell every unit of his endowment at the monopoly price, leaving only residual demand to the larger seller.

The focus of the present paper is on the exercise of strategic power of the single buyer via her purchase decisions.[1] By not purchasing in some period, the buyer can destroy a unit of demand, thereby potentially improving her strategic position, since competition between the sellers intensifies in subsequent periods. Suppose that the purchase decisions of the buyer are perfectly observed by both sellers (the prices they quote may or may not be publicly observed). We find that in any equilibrium, the buyer will be able to obtain a total payoff which is no less than that obtained by purchasing the entire inventory of the smaller seller at zero price. The exercise of strategic buying power has a differential impact upon the sellers when their endowments differ. In contrast with the non-strategic buyer case, we find that the smaller seller loses out, and very often will only get zero profits in equilibrium, while the larger seller makes positive profits. This asymmetric effect arises since the buyer realizes that buying from the smaller seller is costly and reduces her future strategic position since it increases the market power of the larger

[1] There has been some recent interest in buyer strategic power in price setting repeated games under perfect monitoring (see Snyder, 1996) as well as imperfect monitoring (Bergemann and Välimäki (1999a) and Compte (2000)). We discuss this literature in greater detail in the concluding section.

seller. On the other hand, purchase from the larger seller does not reduce future competition equivalently. Our analysis also highlights the distinction between the case when the inventory of the larger of the two sellers is large in relation to total demand ($\omega_A \geq n$), and the case where it is small ($\omega_A < n$). In the large seller case, equilibrium payoffs are unique and efficient, and the total surplus is shared between the buyer and the large seller, with the small seller getting a payoff of 0. In the small seller case, there is always a multiplicity of equilibrium payoffs, and there always exist efficient as well as inefficient equilibria. The efficiency gain relative to the maximally inefficient equilibrium can also be shared between the three parties in a wide variety of ways.

The second part of this paper examines buyer strategic power in a context where transactions are private, so that buyer's purchase decision is not common knowledge between the sellers. We show that there is a dramatic reduction in the strategic power of the buyer, and very often, the strategic buyer can do no better than a myopic one. In particular, in the small seller case, we show that there exist equilibria where the buyer gets zero payoffs.

2 Public Transactions

Let the sellers have endowment $\omega = (\omega_A, \omega_B)$, and assume that there are n time periods where the buyer demands at most one unit in every period. Assume that $\omega_A \geq \omega_B$, and that $\omega_A \leq n$. [2] Index time backwards so that the last period is period 1, and the first period is period n. In each period t, each seller with positive inventory simultaneously quotes the price for a single unit, p_i^t, and the buyer makes a choice to buy from one or none of the sellers. Let $d^t \in \{A, B, \emptyset\}$ denote the buyer's purchase decision, where $\emptyset$ denotes the choice of not buying. We shall assume that prices are privately quoted, so that seller j's price is not observed by seller $i \neq j$. However, in the present section we shall assume the buyer's purchase decision, d^t, is publicly observed by both sellers. The private history of any seller consists at date t, h_i^t, consists of the sequence of own prices and buyers' decisions in the past, $(p_i^\tau, d^\tau)_{n \geq \tau \geq t-1}$, while the private history of the buyer, h_C^t, is the sequence $(p_A^\tau, p_B^\tau, d^\tau)_{n \geq \tau \geq t-1}$. Let $H_i^t, i \in \{A, B, C\}$ denote the set of private histories for a player, and let H_i^n be a singleton set. A pure strategy for seller i a

[2](If $\omega_i > n$, the difference $\omega_i - n$ is irrelevant, and so we simply redefine the endowments, given free disposal, so that $\omega_i \leq n$ for any i.)

sequence of functions $(s_i^t)_{t=n}^1$, where $s_i^t : H_i^t \to \mathbf{R}$ and a pure strategy for a buyer is a sequence of functions $(s_C^t)_{t=n}^1$, where $s_C^t : H_i^t \times \mathbf{R}^2 \to \{A, B, \emptyset\}$.

It is important to note that we have allowed for private strategies, i.e. each player is allowed to condition her behavior on events which are not common knowledge between the players. Our focus is on sequential equilibria. Furthermore, we shall restrict attention to equilibria where each seller i offers a price such that he weakly prefers that the buyer buys from i rather than not buying. Such equilibria are called *cautious* (see Bergemann and Välimäki, 1996) and reflect considerations of trembling hand perfection.[3]

Two examples will make the basic intuition clear. Consider first the case where $n = 2$ and $\omega = (2, 1)$. In this example, $\omega_A \geq n$, i.e. we have the *large seller case,* i.e. the large seller is large relative to market demand.If the buyer buys from seller A at the initial date ($t = 2$), or if she does not buy, this ensures Bertrand competition at the final date and a price of 0, and a continuation value for the buyer of 1. Hence the buyer's payoff in any equilibrium is at least 1. On the other hand, if she buys from seller B at $t = 2$, her continuation value in period 1 is zero, since seller A will be a monopolist at $t = 1$. In consequence, the buyer will buy from B at $t = 2$ only if the price is less than or equal to 0. We conclude therefore that if the buyer buys from either seller at $t = 2$, her payoff will be 1, and seller A's payoff will also be 1. Nor can there be any equilibrium where the buyer fails to buy at $t = 2$ — in this case, seller A's payoff would be zero, and by choosing a price $1 - \varepsilon$ at $t = 2$, A can ensure that the buyer buys. We conclude that the unique cautious equilibrium payoff is $(1, 0, 1)$, and the equilibrium is efficient.[4] Furthermore, the larger seller earns positive payoffs, while the smaller seller makes zero.

Consider next the case where $\omega = (1, 1)$ and $n = 2$. We call this the *small seller case*, since neither seller has sufficient endowment to meet the entire demand by himself. Suppose that the buyer chooses not to buy at the initial date ($t = 2$). This implies that there is Bertrand competition in the final period, and hence the buyer gets the product at price 0. It follows, that in

[3]In an equilibrium which is not cautious, a seller could offer a negative price which makes negative profits, since the buyer buys with probability one from the other seller. Such an equilibrium is not robust to trembles on the part of the buyer.

[4]There exist equilibria which are not cautious where the buyer makes a payoff larger than 1: seller B charges a price $p_B^2 \in [-1, 0)$, A charges $p_A^2 = p_B^2 + 1$, and the buyer buys from A with probability one. If the buyer "trembles" and chooses $d^2 = B$ by mistake, B would earn negative payoffs, and hence such equilibria are not cautious.

any equilibrium, the buyer's utility is at least 1. Indeed, there are three pure strategy equilibria, yielding payoffs $(0,0,1),(1,0,1)$ and $(0,1,1)$ respectively. In the first equilibrium, both sellers charge a strictly positive price and the buyer fails to buy in the initial period, and buys at price zero in the final period. In the second equilibrium, the buyer buys from seller B at price 0 in the initial period, and from seller A at price 1 in the final period, while the third equilibrium reverses the role of the two sellers. Note that the first equilibrium is not efficient — to ensure efficiency, one seller, say seller A has to deviate and offer a price of 0 at $t=2$. This deviation raises the payoff of the other seller, but does not raise either the buyer's or the deviating seller's payoff.

The difference between these two examples is as follows. In the first case, $\omega_A = n$, and the larger seller could meet the entire demand. In consequence, efficiency does not require any coordination with the smaller seller, B. In particular, the buyer could buy from the larger seller for the first $n-\omega_B$ periods, and from either seller in the last ω_B periods. In the second example, in order to ensure efficiency, one seller has to make the sale in the initial period. However, by doing so, this raises the payoff for the other seller, and hence coordination is required.

2.1 The Large Seller Case

We now consider the case where the large seller is large (where $\omega_A \geq n$) more generally and show that equilibrium payoffs are unique.

Proposition 1 *Suppose that $\omega_A \geq n$ and $\omega_B \leq \omega_A$. The payoff vector $(n-\omega_B, 0, \omega_B)$ is the unique cautious equilibrium payoff.*

Proof. Let Γ^t be the class of t period games such that: $\omega_A(t) \geq t$, $\omega_B(t) \leq t$. Let $G^t \in \Gamma^t$ be any game in this class, and let $P(t)$ be the proposition:

$P(t)$: If $G^t \in \Gamma^t$, every cautious correlated equilibrium payoff vector in G^t equals the vector $(t-\omega_B(t), 0, \omega_B(t))$.

We now show that $P(t)$ is true for all t, by induction.

Suppose $t=1$. If $\omega=(1,0)$, correlated equilibrium payoffs are $(1,0,0)$, while if $\omega=(1,1)$, payoff must be $(0,0,1)$. Hence $P(1)$ is true.

We now show that if $P(t-1)$ is true, then $P(t)$ is true.

Let G^t be an arbitrary game in Γ^t. Consider first the case where $\omega_B(t)=t$: Bertrand competition ensures that equilibrium payoffs are $(0,0,\ \omega_B(t))$ so that $P(t)$ is true. Consider next the case where $\omega_B(t) = 0$; clearly, A gets all the payoffs, so $P(t)$ is true in this case as well. Consider finally the case where $0 < \omega_B(t) < t$. If $d^t \in \{ A, \emptyset\}$, $P(t-1)$ implies that the buyer's continuation payoff is $\omega_B(t)$, while if $d^t = B$, $P(t-1)$ implies that her future payoff is $\omega_B(t) - 1$. Since the buyer loses one unit of future payoff by buying from seller B, as compared to not buying, she will only buy from seller B if $p_B^t \leq 0$, and the buyer will only sell if $p_B^t \geq 0$. Hence in any equilibrium where the buyer buys from B, his current payoff is 1, and since the current payoffs of both sellers are zero, we may use $P(t-1)$ to verify that $P(t)$ is true. In any equilibrium where $d^t = A$, the seller must extract all the surplus relative to the buyers alternatives, and hence $p_A^t = \min\{1+p_B^t, 1\}$. Note that $P(t-1)$ implies that the B's continuation value is 0 in all contingencies, and therefore cautiousness implies that his current price $p_B^t \geq 0$. Hence the buyer pays 1 if she buys from seller A in any cautious equilibrium, and by using $P(t-1)$ to compute the total payoffs of all players, we see that $P(t)$ is true. Finally, there cannot be an equilibrium where $d^t = \emptyset$ with positive probability: neither seller A's continuation payoff nor the buyer's continuation payoff varies between the events $d^t = \emptyset$ and $d^t = A$. Hence A can ensure purchase with probability one by offering a price of $1 - \varepsilon$, where ε is arbitrarily small. ■

A corollary of the above proposition is that equilibrium is always efficient in the large seller case.

2.2 The Small Seller Case

Consider the situation where $n > \omega_A \geq \omega_B > 0$, where the larger of the two sellers, A, is small relative to demand — we dub this the small seller case. Clearly, efficiency can only be ensured if some purchases are made from the smaller seller B. On the other hand, the buyer's strategic power implies that he can ensure himself a payoff of at least ω_B. In consequence, efficiency requires sufficient coordination between the players. We shall see that there exists an efficient equilibrium. However, there also exist inefficient equilibria, which reflect the lack of coordination between the players. Furthermore, the division of payoffs amongst the three players is not uniquely determined even if one restricts attention to efficient equilibria. However, the following lemma shows that both the buyer and the larger seller must earn a certain minimum payoff in any equilibrium.

Lemma 2 *If $\omega_B \leq \omega_A$, in any equilibrium the buyer's payoff is at least ω_B and seller's A's payoff is at least $(\omega_A - \omega_B)$.*

Proof. By refusing to buy until period ω_B, the seller ensures that there is Bertrand competition in the last ω_B periods, with the price at 0, ensuring payoff ω_B. On the other hand, by refusing to sell (charging a price greater than 1) in the first $n - \omega_A$ periods, A ensures that the resulting game is such that proposition 1 applies thus ensuring the payoff of at least $\omega_A - \omega_B$. ■

Consider the example where $n = 2$ and $\omega = (1, 1)$. We have already demonstrated that pure strategy equilibrium payoffs include $(0, 0, 1), (1, 0, 1)$ and $(0, 1, 1)$. We now show that any convex combination of these points is a mixed equilibrium payoff. Let the buyer's mixed strategy be denoted by D^t , i.e. it denotes the probability distribution over the set $\{A, B, \emptyset\}$.Let $\lambda_A, \lambda_B, (1 - \lambda_A - \lambda_B) \in (0, 1)$ and consider the class of equilibria where:

At $t = 2 : p_A^2 = p_B^2 = 0$

$D^2 = (\lambda_A, \lambda_B, 1 - \lambda_A - \lambda_B)$ if $p_A^2 = p_B^2 = 0$, $D^2 = (0, 0, 1)$ otherwise. [5]

Such an equilibrium generates payoffs $(\lambda_B, \lambda_A, 1)$ for the three agents.

An intriguing fact is that the buyer's payoff is exactly 1 in every equilibrium that we have considered. Notice that the buyer exercises considerable power, since her decision effectively determines which (if any) of the two sellers gets a payoff of one in the next period. Nevertheless this power does not seem to translate into an increase in the buyer's own payoff. Indeed, one can show that the buyer will never get a payoff greater than 1 in any equilibrium. To get a payoff greater than 1, some seller must price at less than 0. Clearly, such pricing is never optimal at $t = 1$. A seller may price less than 0 at $t = 2$, since such pricing may earn (random) rewards in the future. For example, consider a candidate equilibrium as in the above class, where the price 0 at $t = 2$ is replaced by $p_A^2 = p_B^2 = -x$, and where $\lambda_A = \lambda_B = \frac{1}{2}$. If $x \leq \frac{1}{2}$, such pricing is a optimal from the sellers' point of view, given the buyer's strategy of punishing any deviations by not buying. However, if one seller deviates and increases, the buyer will respond by buying from the other seller, and hence the buyer's punishments are not credible. Hence we conclude that the buyer cannot get a payoff greater than 1 in any equilibrium.

Consider next the example where $n = 3$ and $\omega = (2, 1)$. Lemma 2 implies that the equilibrium payoff vector must weakly greater than $(1, 0, 1)$, i.e.

[5] We omit details of period one behavior, since these are obvious — there is monopoly pricing if only one seller is in the market, and Bertrand pricing otherwise.

both seller A and the buyer are assured of one unit of payoff. We now show that any feasible payoff vector greater than $(1, 0, 1)$ is an equilibrium payoff.

For efficiency, sales must take place in every period, which implies that the seller who makes a sale in the last period must be a monopolist and will therefore earn a payoff of one in this period. Hence in any efficient equilibrium, the seller must make a purchase at price zero at an earlier date. One such equilibrium is where the seller buys from B at price zero at the initial date, and buys at price 1 from A at the subsequent dates, giving rise to the payoff vector $(2, 0, 1)$. In this equilibrium, the buyer will not buy from B at any price greater than 0 (since his continuation payoff is zero), and B is willing to sell at 0 since proposition 1 implies that his continuation value is zero if the buyer does not buy.

Consider next an alternative class of equilibria where the buyer buys from A at $t = 3$ at some price $p_A^3 = x \in [0, 1]$. The buyer buys as long the price is less than or equal to one. Recall that we are assuming that the price quoted by seller A cannot be observed by seller B, and hence B's price at $t = 2$ can only condition upon the purchase decision. Let $p_B^2 = 0$ if $d^3 = A$. The buyer buys from B at price 0 as long as seller A has not deviated in the initial period, i.e. as long as $p_A^3 \leq x$. If $p_A^3 > x$, the buyer chooses not to buy. This ensures that if A deviates at $t = 3$, he loses one unit of payoff in the final period, thus ensuring that he chooses $p_A^3 = x$. This class of equilibria generate payoffs $(1 + x, 0, 2 - x)$ for any $x \in [0, 1]$.[6]

Finally, we show that there is an equilibrium with payoff $(1, 1, 1)$, where the buyer purchases from A at price 1 at $t = 3$, with continuation payoffs $(0, 1, 1)$ for the three players. These are generated by A pricing at 0 at $t = 2$, and B earning monopoly profits in the final period.

[6] If one believes that players will only play equilibria which are Pareto-efficient, then the privateness of prices ensures that seller A will get all the surplus, so that only the payoff $(2, 0, 1)$ is an equilibrium payoff. Since B will price at 0 at $t = 2$ independent of, A can ensure himself a payoff of 1 in the continuation game by pricing higher than 1. Hence no punishment will be effective, unless prices are publicly observed. If prices are publicly observed, then in the subgame that remains (where $n = 2$ and $\omega = (1, 1)$), one can switch between the equilibrium payoffs $(1, 0, 1)$ and $(0, 1, 1)$ depending upon whether A deviates or not at $t = 3$.

2.3 Differentiated Products

Our arguments generalize even when the products are differentiated so that their values to the buyer differ. Assume that the buyer consumes at most one good in each period. Normalize the valuation of the good of seller A to 1, and let the valuation of the good of seller be $1 + \Delta$, where Δ can be positive or negative, but $\Delta > -1$. Proposition 1 generalizes as follows:

Proposition 3 *Suppose that* $\omega_A \geq n$ *and* $\omega_B \leq \omega_A$. *If* $\Delta > 0$, *the payoff vector* $(n - \omega_B, \Delta\omega_B, \omega_B)$ *is the unique cautious equilibrium payoff. If* $-1 < \Delta < 0$, *the payoff vector* $(n - (1 - \Delta)\omega_B, 0, \omega_B(1 - \Delta))$ *is the unique cautious equilibrium payoff.*

Proof. The proof mimics that of proposition 1, and is hence omitted. ■

2.4 Strategic vs Non-Strategic Buyers

Our results may be compared with those obtained by Dudey (1992) for the case of a sequence of buyers (or equivalently, a single myopic or impatient buyer). Dudey shows that if $\omega_B < n \leq \omega_A + \omega_B$, and $\omega_B < \omega_A$, equilibrium payoffs are unique and equal $(n - \omega_B, \omega_B, 0)$. In particular, efficiency is ensured, the buyer gets no surplus, and the smaller of the two sellers gets his monopoly payoff. Furthermore, small differences between the two sellers translate into large payoff differences — if $n = \omega_A = 100$ and $\omega_B = 99$, the equilibrium payoff vector is $(1, 99, 0)$. In contrast, with a strategic buyer, we find that the buyer gains at the expense of the smaller seller. Equilibrium payoffs are unique and efficient only in the large seller case (where the small seller's contribution to social value is zero), and may well be inefficient in the small seller case. Finally, small differences between sellers have only payoff implications of the same order of magnitude — in the example just discussed, the equilibrium payoff vector would be $(1, 0, 99)$, so that the one unit difference in endowment translates into one unit difference in payoffs.

Our analysis also extends to the case where both strategic and non-strategic buyers operate. Suppose that $\omega_A > n$ and $\omega_B < n$, and consider what happens when in the initial period, $n + 1$, a non-strategic buyer enters the market. If the myopic buyer buys from seller B, proposition 1 implies that this raises the continuation value of seller A by one unit. Hence seller A will only be willing to sell to the myopic buyer at a price of 1, and at such a price, A is indifferent between making a sale or not. In equilibrium, the myopic buyer will buy from B at a price of 1, and the presence of such a myopic buyer reduces the payoff of the strategic buyer by 1.

Our results may also be usefully compared with the static case where the market operates only at one date, but where the buyer has n units of demand, and $n \geq \omega_A \geq \omega_B$. If the sellers can offer non-linear prices, each seller can capture his marginal contribution to the buyer's utility, i.e. seller i's payoff is $\min\{(n - \omega_j), \omega_i\}$. Hence, if $\omega_A + \omega_B \leq n$, payoffs are $(\omega_A, \omega_B, 0)$, while if $\omega_A + \omega_B > n$, payoffs are $(n - \omega_B, n - \omega_A, \omega_A + \omega_B - n)$. Suppose that the buyer can strategically choose her level of demand. I.e. given n, she can commit to a level of demand $n' \leq n$, and refuse to buy any additional units. In this case, by choosing n' such that $\omega_A \geq n' \geq \omega_B$, she can ensure herself a payoff of ω_B in the consequent pricing game. Of course, this commitment is not credible in the one-shot context. Our analysis shows that the dynamic context gives such commitments credibility.

3 Private Transactions

We now assume that the buyer's purchase decision is privately observed, i.e. it is perfectly observed only by the seller from whom she makes her purchase, and only imperfectly observed by a seller from whom a purchase has not been made. If the buyer makes a purchase from seller i (i.e. $d = i$), seller i knows that $d = i$, while seller j only knows hat $d \neq j$, i.e. she can infer that $d \in \{i, \emptyset\}$. In the latter event, seller j may obtain a (private) signal which is imperfectly informative. Let Ω denote the finite set of signals which may be observed by seller i when $d \in \{j, \emptyset\}$. Assume that any $\omega \in \Omega$ has strictly positive probability when $d = \emptyset$ and also when $d = j$. In particular, this allows for the possibility that monitoring may be arbitrarily close to perfect.

We now show that imperfect observability implies a significant loss of strategic power for the buyer. This is clearest in the small seller case — there is always an equilibrium where the sellers get all the surplus, and in some cases, there may be no equilibrium where the buyer gets a positive payoff. We turn to an analysis of the small seller case before proceeding to the large seller case.

3.1 The Small Seller Case

Consider the example where $n = 2$ and $\omega = (1, 1)$. In the case of observability of purchase decisions, we saw that the buyer's utility was 1 in any equilibrium. We show first that there exists an equilibrium where the buyer gets utility 0, while the two sellers each get utility 1. Equilibrium strategies are as follows:

At $t = 2$:

$p_A^2 = p_B^2 = 1$.

$d^2 = A$ if $p_A^2 = p_B^2 \leq 1$. $d^2 = i$ if $p_i^2 < p_j^2$ and $p_i^2 \leq 1$.

At $t = 1$:

$p_i^1 = 1$ if $d^2 \neq i$,for $i = A, B$.

$d^1 = A$ if $p_A^1 = p_B^1 \leq 1$. $d^1 = i$ if $p_i^1 < p_j^1$ and $p_i^1 \leq 1$.

This equilibrium has the outcome where the buyer buys from A at price 1 at $t = 2$ and from B at price 1 at $t = 1$. Since each seller makes his maximal feasible profit, clearly neither has any incentive to deviate along the equilibrium path. So consider deviations by the buyer at $t = 2$. If the buyer deviates to $d^2 = \emptyset$, then seller A knows that there has been a deviation, but seller B does not know that there has been a deviation. Hence B continues with his equilibrium strategy, and prices at 1 at $t = 1$. Seller A does not know whether the buyer has deviated to $\emptyset$ or B; however, irrespective of his beliefs, he knows that he can ensure that the buyer purchases with probability one as long as he prices strictly below one, and the tie breaking rule embodied in the buyer's continuation strategy implies this is also the case if $p_A^1 = 1$, regardless of the form of the buyer's deviation. Hence it is optimal for A to price at 1, and the buyer's deviation is unprofitable. Similarly, it is easy to verify that deviating by buying from B at $t = 2$ is unprofitable.

This may be generalized as follows

Proposition 4 *Suppose that* $\omega_A + \omega_B = n$, *and suppose that transactions are not commonly observed. There exists an equilibrium with payoffs* $(\omega_A, \omega_B, 0)$.

Proof. In every period where a seller has positive stock, he chooses the price 1, independent of events in previous periods. The buyer buys from seller A as long as seller A has a positive stock, and as long as $p_A \leq \min\{p_B, 1\}$.Once seller A's stock is exhausted, the buyer buys from seller B, as long as $p_B \leq 1$.

This equilibrium has the path where the buyer buys from A in the first ω_A periods, and from B in the last ω_B periods. If the buyer deviates, say by buying from B when he should be buying from A, it is clear that this deviation does not affect either seller's optimal pricing decision in future periods. By Bayes rule, seller A continues to believe that the buyer will always buy from him. If the buyer does not buy in some period before ω_B, then he will not buy from B in period ω_B, and B will realize that the buyer has deviated at some time in the past. However, this has no implications for further deviations, i.e. the buyer will believe that he will not observe further deviations from the path. By the same logic, if B observes m deviations ($m < \omega_B$), he will believe that there will not be any further deviations. ■

Proposition 5 *Suppose that* $\omega = (1, 1)$ *and* $n = 2$, *and suppose that transactions are private. The payoff* $(1, 1, 0)$ *is the unique equilibrium payoff.*

Proof. Consider first an equilibrium where the buyer buys with probability one at $t = 2$. Fix any such equilibrium where $d^2 = j$ with positive probability along the equilibrium path, and assume that seller i has chosen his equilibrium price p_i^2; then $d^2 \neq i \Rightarrow i$ believes that $d^2 = j$ for any signal that he receives. Hence i will choose the price 1 at $t = 1$ if j does not buy from him at $t = 2$. We show that this implies that $p_j^2 = 1$. If this is not the case, and $p_j^2 < 1$, then j can increase his payoff by choosing $p' \in (p_j^2, 1)$. If the buyer's equilibrium response to this deviation is to choose $d^2 = i$, then j will be a monopolist at $t = 1$, and hence this deviation is beneficial for j. Suppose that the buyer's equilibrium response to j's deviation is to choose $d^2 = \emptyset$. We have established that $d^2 \neq i \Rightarrow i$ believes that $d^2 = j$ for any signal that he receives, and hence i believes that he is a monopolist at $t = 1$, and will choose price 1. Since j can ensure that the buyer buys from him at $t = 1$ by choosing any price $p_j^1 < 1$, equilibrium requires that he price at 1 and the buyer buy from him, and in this case as well the deviation is profitable for j. We conclude that in any equilibrium where the buyer buys with probability one at $t = 2$, he pays a price of 1, and he also buys with probability one at $t = 1$, also at a price of 1.

Consider next a candidate equilibrium where the buyer fails to buy with probability one at $t = 2$. Hence the price of both firms at $t = 1$ equals zero. Suppose now that A offers a price $p_A^1 < 1$. The buyer will certainly buy, since this gives him positive utility and does not affect his continuation value, since seller B cannot observe this deviation. Hence there cannot be an equilibrium where the buyer fails to buy with probability one at $t = 2$.

Finally, we consider the class of candidate equilibria where the buyer randomizes between buying and not buying at $t = 2$. Consider first an equilibrium where $d^2 = \emptyset$ with probability θ and $d^2 = A$ with probability $1 - \theta$, and where A's price at $t = 2$ is p_A^2. Write $V_i^1(d^2 = x)$ for the expected continuation value of agent i ($i \in (A, B, C)$) conditional on the buyer's decision $d^2 = x$ ($x \in \{A, B, \emptyset\}$). Since the buyer must be indifferent between buying and not buying, we must have

$$1 - p_A^2 = V_C^1(d^2 = \emptyset) - V_C^1(d^2 = A) \tag{1}$$

Furthermore, if A charges any price less than p_A^2, the buyer will strictly prefer to buy. Hence A must also be indifferent between making a sale in period two at price p_A^2 and making a sale at $t = 1$ in competition with seller B, i.e.

$$p_A^2 = V_A^1(d^2 = \emptyset) \tag{2}$$

Adding these expressions we obtain

$$V_C^1(d^2 = \emptyset) + V_A^1(d^2 = \emptyset) - V_C^1(d^2 = A) = 1 \tag{3}$$

However, since the total available value at $t = 1$ is 1, this implies that $V_C^1(d^2 = A) = 0$ (and also $V_B^1(d^2 = \emptyset) = 0$). However $V_C^1(d^2 = A) = 0$ implies $p_B^1 = 1$. Lemma 7 in the appendix shows that $p_B^1 = 1$ is inconsistent with $\theta > 0$, and hence we cannot have such an equilibrium where the buyer randomizes between $d^2 = \emptyset$ and $d^2 = A$.

Finally we consider an equilibrium where the buyer randomizes between $d^2 = \emptyset$, $d^2 = A$ and $d^2 = B$. In this case, in addition to the above expressions, one similarly also obtains

$$V_C^1(d^2 = \emptyset) + V_B^1(d^2 = \emptyset) - V_C^1(d^2 = B) = 1 \tag{4}$$

which implies that $V_C^1(d^2 = \emptyset) = 1$, so that at least one seller's price must be zero at $t = 1$ if the buyer does not buy from this seller. However we also have $V_C^1(d^2 = A) = 0$ and $V_C^1(d^2 = B) = 0$, which is inconsistent with this, and hence we cannot have an equilibrium where the buyer randomizes between all three decisions. ■

This proposition shows the sharp discontinuity between imperfect monitoring of transactions, and perfect monitoring. This relates to the literature on imperfectly observed commitments, following Bagwell (1995).[7] Most per-

[7] See also van Damme and Hurkens (1997) and Maggi (2000). Bhaskar and van Damme (1997) analyze a related question in the context of a repeated game with private monitoring.

tinent is the work of Güth, Kirchsteiger and Ritzberger (1998), who show that in any finite game with perfectly observed commitment, there always exists a subgame perfect ("Stackelberg") equilibrium the outcome of which can be approximated under imperfect observability with small noise. The above proposition shows that this is not the case in our model of strategic pricing. If $n = 2$ and $\omega = (1, 1)$,with perfect observability of purchases, the buyer gets a payoff of 1 in any equilibrium, whereas with imperfect observability, she gets a payoff of 0 in any equilibrium. Of course the proposition of Güth et. al. does not apply in the present context since strategy sets are infinite in our pricing game.

3.2 The Large Seller Case

The discussion of the small seller case suggests that with private transactions, the buyer completely loses his strategic power. Indeed, she does no better than the myopic buyer does. We now show that this is not true in the large seller case. For example, let $n = 2$ and $\omega = (2, 1)$. The following equilibrium gives payoffs $(1, 0, 1)$:

At $t = 2 : p_A^2 = 1, p_B^2 = 0$

$d^2 = A$ if $p_A^2 \leq \min\{1, 1 + p_B^2\}$

$d^2 = B$ if $p_B^2 < p_A^2 - 1$ and $p_B^2 \leq 1$.

$d^2 = \emptyset$ if $p_B^2 > 1$ and $p_A^2 > 1$.

At $t = 1 : p_A^1 = 0$ if $d^2 = A$, $p_A^1 = 1$ if $d^2 \neq A$.

$p_B^1 = 0$ if $d^2 \neq B$.

The crucial point in the large seller case is that if that the small seller knows that the large seller always has positive inventory in any period, independent of past events. Consequently, provided that the buyer buys from the large seller A, A knows that she has not bought from B, and therefore Bertrand competition results.

Proposition 6 *Suppose that $\omega_A \geq n$ and $\omega_B \leq \omega_A$. There exists an equilibrium with payoff $(n - \omega_B, 0, \omega_B)$.*

Proof. In the first $n - \omega_B$ periods, seller A prices at 1 and seller B prices at 0, and the buyer buys from seller A, as long as $p_A^t \leq \min\{1, 1 + p_B^t\}$. If A deviates by pricing higher than 1, the buyer buys from B. If there are no deviations, in the last ω_B periods, both sellers price at 0. If the buyer does not buy from A for some k periods in the first $n - \omega_B$ periods, then the sellers price at 0 only in the last $\omega_B - k$ periods. ■

4 Concluding Comments

There has been some recent interest in the power of a strategic buyer. Snyder (1996) considers a repeated procurement auction which is a repeated game with perfect monitoring. Bergemann and Välimäki (1999a) and Compte (2000) consider repeated pricing games with imperfect and private monitoring. The present paper differs from this literature since our focus is on a dynamic game rather than a repeated game, where the inventory levels of the sellers are directly payoff relevant. Our focus is on the lack of observability of transactions, whereby the inventory levels are no longer common knowledge between the sellers as time proceeds.

Our paper is also related to the question of efficiency in dynamic games of price competition. Work in this area includes Bergemann and Välimäki (1996, 1999b), and Felli and Harris (1996). Our main finding is that strategic buying power and perfect observability may result in inefficiencies.

5 Appendix: Bertrand-Edgeworth with Random Capacities

The proof of proposition 5 requires the following lemma pertaining to a one period model of Bertrand competition with random endowments. Suppose that the seller has demand 1, and suppose that $\omega_B = 1$ with probability one. Nature selects $\omega_A = 0$ with probability θ_A and $\omega_A = 1$ with probability $1 - \theta_A$, and each seller observes his own realized endowment but does not observe the endowment of his rival.

Lemma 7 *In any equilibrium of the above Bertrand game with random endowments, seller B's price is less than* 1 *with positive probability.*

Proof. Suppose that $p_B = 1$ with probability one. If $\omega_A = 1$, then A can ensure himself of a payoff arbitrarily close to 1 by choosing a price $1 - \varepsilon$. Hence equilibrium requires that A also choose a price of 1, and that the buyer buys from A when $p_A = p_B = 1$. However, this implies that B can ensure himself a profit arbitrarily close to 1 by choosing a price arbitrarily close to 1, whereas he earns only θ_A by choosing a price of 1. Hence there cannot be an equilibrium where $p_B = 1$ with probability one. ■

More generally, we may consider the following one period model of Bertrand competition with random endowments. Suppose that for each $i \in \{A, B\}, \omega_i =$

0 with probability θ_i and $\omega_i = 1$ with probability $1 - \theta_i$, and suppose that the seller had demand 1. Without loss of generality, we may index players so that $\theta_B \geq \theta_A$. Assume also that $\theta_B > 0$, since otherwise one has standard Bertrand competition.. We claim that the following mixed strategies constitute an equilibrium:

B chooses a price according to the uniform distribution on $[\theta_B, 1)$.

A chooses price 1 with probability mass $\theta_A - \theta_B$; the remaining probability is distributed uniformly on prices between $[\theta_B, 1)$.

The buyer chooses the seller with the lower price if prices differ. If prices are equal, she chooses either seller with equal probability.

We verify that these strategies constitute an equilibrium. If any player i chooses the lowest price θ_B, he sells with probability one, and earns profits θ_B. Similarly, if A chooses price 1, he sells with probability θ_B, and earns θ_B in expected payoff. Let $F_j(p)$ be the distribution function of j's mixed strategy. The payoff of player i from price $p \in (\theta_B, 1)$ is given by

$$U_i(p) = \theta_B p + (1 - \theta_B) F_j(p) p \tag{5}$$

Differentiating, we get

$$\frac{\partial U_i}{\partial p_i} = \theta_B + (1 - \theta_B)[p F_j'(p) + F_j(p)] = 0 \tag{6}$$

Which is satisfied if F_j is given by the uniform distribution on $(\theta_B, 1)$.

Note that the utility of the buyer is less than $(1 - \theta_B)$, since prices are greater than θ_B.

References

[1] K. Bagwell (1995). "Commitment and Observability in Games", *Games and Economic Behavior,* **8**, 271-280

[2] Bergemann, D., and J. Välimäki (1996). "Learning and Strategic Pricing", *Econometrica* **64**, 1125-1149.

[3] Bergemann, D., and J. Välimäki (1999a). "Dynamic Common Agency", mimeo.

[4] Bergemann, D., and J. Välimäki (1999b). "Strategic Buyers and Privately Observed Prices", mimeo

[5] Bhaskar, V. and E. van Damme, (1997). "Moral Hazard and Private Monitoring", mimeo.

[6] Compte, O. (2000). "Secret Price Cutting and Strategic Buyers", mimeo

[7] Cooper, D.J., and R. G. Donaldson. (1998). "A Strategic Analysis of Corners and Squeezes", *Journal of Financial and Quantitative Analysis* **33,** 117-137.

[8] E. van Damme, and S. Hurkens, (1997). "Games with Imperfectly Observable Commitment", *Games and Economic Behavior* **21**, 282-308.

[9] Dudey, M. (1992). "Dynamic Edgeworth-Bertrand Competition", *Quarterly Journal of Economics* **6,** 1461-1477.

[10] Felli, L, and C. Harris, (1996). "Job Matching, Learning and the Distribution of Surplus", *Journal of Political Economy.*

[11] Güth, W., G. Kirchsteiger and K. Ritzberger, (1998). "Imperfectly Observable Commitments in n-Player Games", *Games and Economic Behavior,* **23,** 54-74.

[12] Maggi, G.,(1999). "The Value of Commitment with Imperfect Observability and Private Information", *Rand Journal of Economics* **30,** 555-574.

[13] Snyder, C. (1996) "A Dynamic Theory of Countervailing Power", *Rand Journal of Economics* **27,** 747-769.