

Dynamic Countervailing Power under Public and Private Monitoring

V. Bhaskar*

Dept. of Economics
University College London
Gower Street
London WC1E 6BT.
Email: v.bhaskar@ucl.ac.uk

June 2013

Abstract

We examine buyer strategic power in the model of dynamic Bertrand-Edgeworth competition. Two sellers with a limited inventory sell to a single buyer, who has a consumption opportunity in each period. 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 publicly observed, we find that that a strategic buyer can do significantly better than non-strategic buyers; strategic power may also 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, and in some cases, may be completely eliminated. This highlights the sharp discontinuity between the equilibrium outcomes between perfect and imperfect monitoring.

Keywords: dynamic games, imperfect monitoring, strategic buyer, Bertrand-Edgeworth competition. JEL codes: D43,D92.

*Thanks to Helmut Bester, Juuso Välimäki and many seminar audiences for helpful comments.

1 Introduction

The power of large buyers such as supermarkets and retailers has drawn the attention of policy-makers (see Clarke et al. [10]). It has also captured the popular imagination and has credibility among consumers – LetsBuyIt.com, an internet based company, advertised extensively based on the idea that large buyers earn substantial discounts. Although buyer power is viewed with concern in markets where sellers are small (e.g. farmers), in oligopolistic markets, this *countervailing power* is viewed favorably. What are the sources of countervailing power, and how is it exercised? Bargaining models, with a monopoly seller and one/several buyers, provide one answer (see Stole and Zweibel [28], Dobson and Waterson [14] and von Ungern-Sternberg [29], for example). Most markets are however oligopolistic, and this may afford the buyer greater strategic possibilities. Indeed, Galbraith, the most influential proponent of the thesis of countervailing power, believed that bilateral monopoly was a less plausible as its foundation. He argued that “oligopoly facilitates the exercise of countervailing power by enabling a strong buyer to play one seller off against another” ([17], p. 127).¹ Anton and Yao [1], [2] model such considerations in a static context. They examine a single buyer who may split his orders between two sellers, and his incentives to do so when the sellers’ cost functions are strictly convex. Inderst [20] considers multiple buyers within a similar set up and shows that the incentive to split orders increases.

In a dynamic context, the idea that a strategic buyer can judiciously un-

¹Indeed, he argues that countervailing power is an important reason why oligopoly is more progressive than pure monopoly. “One can hardly doubt that, in general, it will be much easier for countervailing power to break into a position of market strength maintained by an imperfect coalition of three, four or a dozen firms than into a position held by one firm. When there is more than one firm in the market there are opportunities for playing one off against another. Mistrust and uncertainty can be developed in the mind of one entrepreneur as to the intentions and good faith of others.” (p. 159).

dermine seller collusion appears in some recent papers. Snyder [26] (see also [27]) studies a repeated procurement auction, where the buyer can defer purchases at some cost. By doing so, the buyer can increase the level of demand, thereby reducing the extent of collusion, as in Rotemberg and Saloner's [25] model of "price wars in booms". Snyder shows that this strategic possibility limits the degree of collusion among sellers. Bergemann and Välimäki [6] and Compte [11] consider the role of a strategic buyer in a context where sellers' ability to collude is impaired by their ability to make secret price cuts. Such a private monitoring environment poses serious problems in sustaining collusion, as the recent literature on repeated games with private monitoring shows. Bergemann-Välimäki and Compte show that a strategic buyer can impose serious limitations upon collusion, since she can simulate a situation where one seller suspects that the other has secretly cut price.²

The focus of the present paper is on the implications of buyer's strategic power in the context of a dynamic, non-repeated game environment, with limited seller capacities. That is, we introduce a strategic buyer in the context of the model of dynamic Bertrand-Edgeworth competition due to Dudev [15].³ Dudev considers a model with a sequence of one-shot buyers, each of whom enters the market for one period only. His essential finding is counter-intuitive — the smaller of the two sellers is often much better off than the larger. For example, if there is a sequence of 100 buyers, each of whom demands one unit, and the two sellers have inventory of 100 and 99 units respectively, then the smaller seller sells 99 units at the monopoly price while the large seller only sells 1 unit. This model also displays a dramatic "discontinuity" — if the smaller buyer were to have one more unit, or if there was one

²Bergemann and Välimäki demonstrate this in the context of a finitely repeated interaction, while Compte studies a simple two period model where sellers are able to communicate and commit to a side-payment mechanism in order to sustain collusion.

³Ghemawat and McGahan [18] apply this model to the context of the large turbine generator industry in the US, while Cooper and Donaldson [12] consider its relevance to financial markets. In the former context at least, strategic buyers seem a real possibility.

less buyer, then one would move from full monopoly pricing to competitive pricing.

Consider now the implications of having a strategic buyer, who seeks to consume one unit of the product in each of 100 periods.⁴ The buyer's consumption opportunities are perishable, which allows her, in a dynamic context, to manipulate the demand-supply configuration in her favor, by forgoing consumption in a period. If the buyer's purchase decisions are public, this possibility — even if not exercised — confers significant countervailing power to the strategic buyer, in comparison with a non-strategic buyer. Thus in any equilibrium, the buyer must get a payoff which corresponds to getting the product at marginal cost for 99 periods. Furthermore, it is the larger seller who earns more than the smaller seller. We show that the analysis depends upon whether the smaller seller is essential or not. If the larger seller can meet the entire consumption requirements of the buyer from his inventory, then we say that the smaller seller is *inessential*. In this case, equilibrium outcomes are unique, with the buyer and the large seller splitting the total surplus, in a ratio that depends upon the degree of competition offered by the smaller seller. However, if the smaller seller is *essential* since total demand cannot be met by the larger seller, then there is a multiplicity of equilibrium payoffs. In either case, we provide a characterization of equilibrium outcomes.

This part of the paper is related to recent work by Anton, Biglaiser and Vettas [3] who consider a two-period model of a durable good, with two sellers and a single buyer. They show that the buyer has an incentive to split his orders between the sellers in initial period, as a way of increasing price competition in the future. They also find that this hurts the buyer. The key difference is that in our paper, there are two strategic decisions – whether

⁴Contrast this with the Rotemberg-Saloner [25] or Snyder [26] repeated game intuition with unbounded capacities, where prices are lower when demand is greater.

to buy or not, and who to buy from, both of which play an important role in maintaining buyer power. We are also able to characterize equilibrium outcomes for an arbitrary finite horizon. Somewhat different strategic effects arise in models where suppliers are subject to learning by doing, as in Cabral and Riordan [9] and Lewis and Yildirim [22], where increased purchase from one seller makes that seller more competitive and vis-a-vis the opponent.

The second part of our analysis relaxes the assumption that transactions are public and are perfectly observed. While the assumption of public transactions may be appropriate in some contexts (e.g. government agencies, supermarkets or retailers), in many other instances it may be more appropriate to assume that the buyer's purchase decisions are private — i.e. the purchase decision is perfectly observed only by the buyer herself and the seller she buys from. We allow for the possibility that a seller not involved in a transaction may observe an imperfect (though possibly very accurate) signal of the buyer's purchase or non-purchase. With private transactions, the buyer's countervailing power is considerably undermined. Indeed, in some contexts, she will be able to do no better than a myopic buyer, in *any* equilibrium.⁵ This result is striking, and surprising. While Bagwell [4] pointed out the implications of imperfect observation of the actions of a Stackelberg leader, subsequent work has shown that the leader's commitment power can be preserved via mixed strategy equilibria — see van Damme and Hurkens [13] and Guth et al. [19]. Maggi [23] also shows that commitment power can be restored if the leader is subject to payoff shocks. In contrast, we show that the buyer's commitment power cannot be preserved even in mixed strategy equilibria. Intuitively, our result arises because we have a price-setting situation, where randomization is impossible to sustain.

The remainder of this paper is organized as follows. Section 2 considers

⁵More precisely, we find that the set of equilibrium outcomes with public transactions are disjoint from the set of equilibrium outcomes with private transactions.

the case of public transactions that are perfectly observed. Section 3 turns to an analysis of private transactions. The final section concludes.

2 Public Transactions

Two sellers, A and B , each have a finite endowment of inventory, $\omega_A, \omega_B \in \mathbb{N}$, where $\omega_A \geq \omega_B$. 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 value the good at zero. Assume that there is free disposal so that we may as well assume $\omega_A \leq n$. 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. By public transactions, we mean that the buyer's purchase decision, d^t , is publicly observed by both sellers, an assumption we shall maintain for the present section. The prices quoted by seller j may either be public, or they may be privately quoted, in which case so that seller j 's price is not observed by seller $i \neq j$. Our results do not differ very much in these two cases, so we consider both cases together.

Consider first the case of private prices. 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 sequence of functions $(s_i^t)_{t=n}^1$, where $s_i^t : H_i^t \rightarrow \mathbf{R}$ and a pure strategy for a buyer is a sequence of functions $(s_C^t)_{t=n}^1$, where $s_C^t : H_C^t \times \mathbf{R}^2 \rightarrow \{A, B, \emptyset\}$. In the case where the prices are public, $H_i^t = H_C^t$ for $i \in \{A, B\}$ so that all strategies can be conditioned upon the public

history.

Our focus is on perfect Bayesian 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, [5]) and reflect considerations of trembling hand perfection.⁶

A useful static benchmark for our analysis is when the market operates only at one date, but where the single buyer has n units of demand, and $n \geq \omega_A \geq \omega_B$. We assume that each seller simultaneously makes a take-it-or-leave-it offers to the buyer. Without loss of generality, we may assume that seller i offers her whole inventory, ω_i (since the seller does not value the good) at price P_i .

Proposition 1 *If the market only operates at the initial date, each seller can capture his marginal contribution to the buyer's utility, and earn $\min\{(n - \omega_j), \omega_i\}$. Thus, if there is no excess supply ($\omega_A + \omega_B \leq n$), each seller sells his endowment at unit price of ω_i , and the buyer gets no surplus. If $\omega_A + \omega_B > n$, so that there is excess supply, seller A's payoff is $n - \omega_B$, seller B earns $n - \omega_A$, and the buyer gets a surplus of $\omega_A + \omega_B - n$.*

Proof. Suppose that $\omega_A + \omega_B \leq n$. Then if seller i charges ω_i , it is optimal for the buyer to accept this offer, independent of P_j . Since it cannot be optimal for the seller to charge less than ω_i (the buyer will still buy if he increases the price slightly), this is the unique equilibrium. Now suppose that $\omega_A + \omega_B > n$. If seller i chooses $P_i = n - \omega_j$, then it is optimal for the buyer to buy, independent of the offer P_j . Any larger value of P_i is unacceptable to the buyer if $P_j = n - \omega_i$, and thus this is the unique equilibrium outcome. ■

⁶In 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.

The static model illustrates the benefit to the buyer from reducing her level of demand. Suppose that the buyer can strategically choose her level of demand — i.e. she can commit to any level of demand $n' \leq n$, and refuse to buy any additional units. In this case, by choosing $n' = \omega_B$, 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 — given the take-it-or-leave-it offers of the sellers, the buyer would have no incentive to carry out her commitment. However, the dynamic context gives such commitments credibility, since consumption opportunities are perishable. By not purchasing in some period, she can destroy a unit of demand, thereby potentially improving her strategic position, since competition between the sellers intensifies in subsequent periods. The buyer always has the option of not consuming until the number of remaining periods equals ω_B , which ensures that she will get the good at a price of zero for all these periods. It follows 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.

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 *inessential seller case*, i.e. the small seller is inessential to meet total 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.⁷ 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 *essential 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 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,

⁷There 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.

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 Inessential Small Seller Case

We now consider the case where the large seller is large enough that the smaller seller is inessential – i.e. $\omega_A \geq n$.

Proposition 2 *Suppose that $\omega_A \geq n$ and $\omega_B \leq \omega_A$ and transactions are public. The payoff vector $(n - \omega_B, 0, \omega_B)$ is the unique 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 $G^t \in \Gamma^t$, the unique cautious 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 $\omega = (1, 0)$, equilibrium payoffs are $(1, 0, 0)$, while if $\omega = (1, 1)$, the 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 is a monopolist in every period and 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 buyer's 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 Essential 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 case where the small seller is essential. 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 first part of the following proposition shows that both the buyer and the larger seller must earn a certain minimum payoff in any equilibrium. Since the maximum available payoff is given by $\min\{\omega, n\}$, and since the first part of the proposition dictates the distribution of ω_A units of payoff,

this implies that we have a *residual* payoff, $\rho = \min\{\omega_B, n - \omega_A\}$, whose distribution has to be determined. The proposition shows that this residual may be distributed to either seller, or it may be destroyed.

Proposition 3 *If $\omega_B \leq \omega_A < n$, in any equilibrium the buyer's payoff is at least ω_B and seller's A 's payoff is at least $(\omega_A - \omega_B)$. There always exist equilibria where a) the residual payoff ρ accrues to seller A , b) ρ accrues to seller B and c) the residual payoff is destroyed since there is no trade in ρ periods. .*

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, A can charge a price of 1 in the first $n - \omega_A$ periods. If the buyer does not buy, then the resulting game is such that proposition 2 applies thus ensuring A the payoff of at least $\omega_A - \omega_B$.

For the second part, we prove part (c) first. In any period t such that the residual payoff is positive, $p_B > 0$ and $p_A = 1$, and the buyer does not buy. At the end of this phase, when the residual payoff is zero, proposition 2 applies, and hence the A gets $(\omega_A - \omega_B)$ while the buyer gets ω_B . To verify that there is no profitable deviation in the first phase, note that the buyer reduces his payoff in phase two by one unit if he buys from B , and hence will not buy from B at any strictly positive price. Seller A also reduces his phase two payoff by one unit if he sells in phase one, and hence will not sell at any price less than 1. Hence there is no profitable trade either between A and the buyer or between B and the buyer in phase 1.

Proof of part a): In the first phase, which lasts as long as $\omega_B(t) > 0$ and $\omega_B \leq \omega_A$, seller B prices at 0, while seller A prices at 1 if $\omega_B(t) < t$ and at 0 if $\omega_B(t) \geq t$. The buyer buys from B provided that the price is not greater than zero. Once B 's stock is exhausted, seller A sells at a price of 1. If B deviates with a higher price, then the buyer does not buy provided that

$\omega_B(t) < t$ and buys from A otherwise.

Proof of part b): For the first $\omega_A - \omega_B$ periods, the buyer buys from A at a price of one, while B chooses a strictly positive price. At this point, the proof of part 1 implies that there exists an equilibrium where the residual payoff ρ remaining goes to B . ■

Although the residual payoff may go to either seller in some equilibrium, can it also accrue to the buyer? We now show this may or may not be possible. 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 $(0, 1, 1)$. Indeed, any convex combination of these points is a mixed equilibrium payoff.⁸ 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 his price, 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.

For an example where the buyer gets the residual payoff, let $n = 3$ and $\omega = (2, 1)$. Let us consider a class of equilibria where the buyer buys from

⁸This may be achieved as follows: both sellers choose a price of zero at $t = 2$, and given that they do so, the buyer randomizes appropriately between all elements of the set $\{A, B, \emptyset\}$. If any seller deviates and chooses a positive price, the buyer chooses \emptyset for sure.

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. 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 generates payoffs $(1 + x, 0, 2 - x)$ for any $x \in [0, 1]$.

The results of this section are also related to the work examining the efficiency of outcomes in dynamic games of price competition, such as Bergemann and Välimäki [5],[6] and Felli and Harris [16]. Whereas the main point of these papers is that decisions are often dynamically efficient, our main finding is that strategic buying power may sometimes result in inefficiencies.

2.3 Extensions

We now consider some extensions, to show how the results generalize.

2.3.1 Discounting

Our assumption, of zero discounting, is a simplifying one, and the results are robust to players being somewhat impatient.⁹ To illustrate this, let us consider the essential small seller case where $\omega = (1, 1)$ and $n = 2$, and let $\delta < 1$ be the common discount factor for all agents. In this case, it is impossible to have a pure strategy equilibrium where the buyer fails to buy with probability one in the initial period. In such an equilibrium, both sellers earn zero payoffs, while the buyer earns δ . But in this case, a seller can offer the product a price slightly below $1 - \delta$ in the initial period, and the buyer will buy. However, one can construct a mixed strategy equilibrium where both sellers choose price $1 - \delta$ in the initial period, $t = 2$. On the equilibrium path, buyer randomizes between all three options at $t = 2$, choosing each

⁹I am grateful to Helmut Bester for raising this question.

buyer with probability $\theta < 0.5$, and not buying with probability $1 - 2\theta$. In the final period, the price equals zero (resp. one) if both (resp. one) sellers are active. If a seller deviates by raising his price at $t = 2$, the buyer chooses not to buy from either seller, making such a deviation unprofitable. The buyer will buy from the deviating seller if he reduces price at $t = 2$. For this to be unprofitable, the equilibrium payoff for a seller must be greater than the supremum of his deviation payoff, i.e. $(1 - \delta)\theta + \delta\theta \geq 1 - \delta$, i.e. $\theta \geq 1 - \delta$. We therefore see that discounting bounds the efficiency, but as $\delta \rightarrow 1$, the probability of the buyer not buying in the initial period can increase to one.

2.3.2 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 2 generalizes as follows:

Proposition 4 *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 2, and is hence omitted. ■

2.3.3 Strategic vs Non-Strategic Buyers

Our results may be compared with those obtained by Dudey [15] 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 2 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.

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. We shall also assume that the prices offered by seller i are not observed by the other seller. If the buyer makes a purchase from seller i (i.e. $d = i$), seller i knows that $d = i$, while seller j only knows that $d \neq j$, 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. Apart from this full support assumption, we do not have to make any further assumption. Let $\Omega_i = \Omega \cup \{i\}$. Thus in any period, the buyer's decision results in signal $\omega \in \Omega_i$, where the signal i is perfectly informative and indicates that the buyer has bought from i .

The private history of any seller consists at date t , h_i^t , consists of the sequence of own prices and signals about the buyers' decisions in the past, $(p_i^\tau, \omega_i^\tau)_{n \geq \tau \geq t-1}$, 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 sequence of functions $(s_i^t)_{t=n}^1$, where $s_i^t : H_i^t \rightarrow \mathbf{R}$ and a pure strategy for a buyer is a sequence of functions $(s_C^t)_{t=n}^1$, where $s_C^t : H_C^t \times \mathbf{R}^2 \rightarrow \{A, B, \emptyset\}$.

Fix a pure strategy profile $s = (s_A, s_B, s_C)$. This profile induces a sequence of prices by the sellers and decisions by the buyer, along the equilibrium path. Let h_i^t be a private history for seller i . We shall say that seller i has an *observable deviation* at this history if either he has deviated by choosing prices that are different from those indicated by s_i , or if the buyer has chosen $d^i \neq i$ when s_C induces $d^i = i$, or has chosen $d^i = i$ when s_C induces $d^i \neq i$. It is noteworthy that our game has a mix of observable deviations and deviations that may statistically detectable (since the seller observes informative signals in Ω), but that are not perfectly observed. This differs from say repeated games with private monitoring, where there are no observable deviations.

Given that our equilibrium concept is perfect Bayesian equilibrium and that the sellers have imperfect information, we will need to specify their beliefs as well as the equilibrium strategies. However, to economize on space, we will not specify beliefs at information sets that are reached with positive

probability along the equilibrium path – the beliefs here must be derived via Bayes’ rule. We will only specify beliefs at seller information sets that arise after an observable deviation – when the buyer is supposed to make a purchase from seller i with probability one in period t , but in fact does not do so.

We now show that imperfect observability implies a significant loss of strategic power for the buyer. This is clearest when the small seller is essential — 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. Since this case has the most striking results, we analyze this first before proceeding to the case where the smaller seller is inessential.

3.1 Essential 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.

The equilibrium strategies are as follows:

At $t = 2$:

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

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

At $t = 1$:

$$p_i^1 = 1 \text{ if } d^2 \neq i, \text{ for any } \omega \in \Omega, \text{ for } i = A, B.$$

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

If the buyer fails to buy from seller A at $t = 1$, A ’s beliefs regarding d^1 can be any probability distribution with support in the set $\{B, \emptyset\}$.

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. Regardless of the signal that B obtains, he continues to believe that he is a monopolist. 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 5 *If $\omega_A + \omega_B = n$ and transactions are private, then 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 $p_B \leq 1$. If seller A has an observable deviation of the form $d^t \neq A$ in any period $t > \omega_B$, then A believes that $d^t = B$ as long as this is feasible, and $d^t = \emptyset$ otherwise. If seller B has an observable deviation where $d^t \neq B$ in any period $t \leq \omega_B$, then B believes that the buyer has always bought from A in all previous periods τ where $d^\tau \neq B$.

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, since seller A continues to believe that the buyer will always buy from him. If the buyer does not buy from B in period ω_B , B continues to believe that he is a monopolist and will not reduce his price. By the same logic, if B observes m deviations ($m < \omega_B$), he will believe that there will not be any further deviations. ■

The above proposition is reminiscent of the work of Bagwell [4], who studied a finite leader-follower game, and showed that the leader loses his commitment power when his action is observed imperfectly by the follower. We now generalize this result in several dimensions. First, we show that the buyer completely loses his power in *any* pure strategy equilibrium. Second, we show that this is also true for mixed strategy equilibria for a particular instance, when there are two periods.

Proposition 6 *Suppose $\omega_A + \omega_B = n$. In any pure strategy equilibrium, the buyer always buys in every period, and pays a price of 1.*

Proof. See appendix. ■

Can the buyer's strategic power be retained in *some* equilibrium of the game with small sellers? While we are not able to provide a complete answer to this question, the following proposition has a striking negative result for a specific example.

Proposition 7 *Suppose that $\omega = (1, 1)$ and $n = 2$, and suppose that transactions are private. The payoff $(1, 1, 0)$ is the unique equilibrium payoff.*

The proof of this proposition 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 8 *In any equilibrium of the above Bertrand game with random endowments, seller B's price is less than 1 with positive probability.*

Proof. See Appendix. ■

We now turn to the proof of proposition 7.

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 the buyer 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 $p_A^2 < 1$. 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

$$1 - p_A^2 = V_C^1(d^2 = \emptyset) - V_C^1(d^2 = A). \quad (1)$$

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.

$$p_A^2 = V_A^1(d^2 = \emptyset). \quad (2)$$

Adding these expressions we obtain

$$V_C^1(d^2 = \emptyset) + V_A^1(d^2 = \emptyset) - V_C^1(d^2 = A) = 1. \quad (3)$$

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 8 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

$$V_C^1(d^2 = \emptyset) + V_B^1(d^2 = \emptyset) - V_C^1(d^2 = B) = 1, \quad (4)$$

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 [4].¹⁰ Most pertinent is the work of Güth, Kirchsteiger and Ritzberger [19], who show that in any finite game with perfectly observed commitment and generic payoffs, 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 every equilibrium.¹¹ Indeed, in some contexts, she will be able to do no better than a myopic buyer, in *any* equilibrium.¹² This result is striking, and surprising. Intuitively, our result arises because in our richer environment with price-setting situation, randomization is impossible to sustain.

¹⁰See also van Damme and Hurkens [13] and Maggi [23]. Bhaskar and van Damme [5] analyze a related question in the context of a repeated game with private monitoring.

¹¹The proposition of Güth et. al. does not apply in the present context since our game is more complex – possibly the most important factor is that strategy sets are infinite in our pricing game. Morgan and Vardy [24] provide an example of a leader-follower game with continuum action sets where a discontinuity arises but this involves rather different considerations, due to the strict convexity of payoff functions. See also Bhaskar [8].

¹²More precisely, we find that the set of equilibrium outcomes with public transactions are disjoint from the set of equilibrium outcomes with private transactions.

3.2 Inessential Small Seller Case

The discussion of the essential 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 when the small seller is inessential, For example, let $n = 2$ and $\omega = (2, 1)$. The following equilibrium gives payoffs $(1, 0, 1)$:

$$\begin{aligned} \text{At } t = 2 : p_A^2 &= 1, p_B^2 = 0 \\ d^2 = A &\text{ if } p_A^2 \leq \min\{1, 1 + p_B^2\} \\ d^2 = B &\text{ if } p_B^2 < p_A^2 - 1 \text{ and } p_B^2 \leq 1. \\ d^2 = \emptyset &\text{ if } p_B^2 > 1 \text{ and } p_A^2 > 1. \\ \text{At } t = 1 : p_A^1 &= 0 \text{ if } d^2 = A, p_A^1 = 1 \text{ if } d^2 \neq A. \\ p_B^1 &= 0 \text{ if } d^2 \neq B. \end{aligned}$$

The crucial point is that when the small seller is essential, 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. This may be generalized as follows.

Proposition 9 *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. ■

The results here, when the small seller is inessential, show that some strategic power can be maintained despite the lack of perfect observability. The key here is that buying from the large seller provides the relevant

commitment device for the buyer, even though not buying is not a credible commitment. This suffices to ensure that the buyer has some bargaining power.

4 Concluding Comments

We have examined how a strategic buyer may use her purchase decisions in order to improve her competitive position vis-a-vis oligopolistic sellers. Our analysis has highlighted the importance of the monitoring structure. If the buyer's decisions publicly and perfectly observed by both sellers, the buyer has considerable power. On the other hand, if purchases are private, and are only imperfectly observed by third-parties, then the buyer may lose her power substantially. Thus a second contribution of this paper is to highlight the role of imperfect observation in the context of a rich economic environment, such as dynamic price competition.

5 Appendix

We present the proof of proposition 6.

First, we show that in any pure strategy equilibrium where the buyer always buys, i.e. $d^t \in \{A, B\} \forall t$, the price he pays must equal 1. We prove this by induction along the equilibrium path.

Consider histories where neither A nor B has an observable deviation, i.e. where the buyer has bought from the buyer he was supposed to buy from at that date. Suppose that $t = 1$ (i.e. it is the terminal period), and that buyer buys from seller i at this history in equilibrium. Seller i believes that he is a monopolist, and hence price, p_i^t must equal 1. Note that this is true as long as i does not have an observable deviation – if, for example, the buyer did

not buy from j at $t = 2$ and instead chose \emptyset , seller i will still price at 1 since i does not have an observable deviation.

Consider a history at $t = 2$ where there is no observable deviation to the seller who sells along the equilibrium path at $t = 2$. If this seller also sells at $t = 1$, then this seller believes that he is a monopolist for both these periods and thus price must equal 1 at $t = 2$. Suppose that the buyer buys from A at $t = 2$ and from B at $t = 1$. Let the price paid to A be p_A^2 . Since B 's continuation payoff in this equilibrium is 1, $p_B^2 \geq 1$. Now A has the option of choosing a price of 1 at $t = 2$, and if C refuses to buy, choosing a price of $1 - \varepsilon$ at $t = 1$, thereby achieving a payoff of $1 - \varepsilon$ for any $\varepsilon > 0$. Thus the only equilibrium price is $p_A^2 = 1$.

We have therefore established that if the seller who sells at date 2 has no observable deviation, then he must price at $t = 1$, and as long as there has been no deviation by anyone, the buyer must buy with probability one. Note that the buyer's continuation value along this equilibrium path is zero, since he buys at a price of 1 at $t = 2$ and at $t = 1$. Since equilibrium requires that the buyer cannot increase his continuation value by not buying at $t = 2$, if the buyer deviates and does not buy from some seller j at $t = 2$, then the price charged by j at $t = 1$ must equal 1. Furthermore, seller j must believe that he will make a sale with probability one at $t = 1$, since otherwise he would have an incentive to reduce his price.

Suppose now that for every $t \in \{1, 2, \dots, k - 1\}$, the equilibrium price equals 1 and the buyer buys with probability one along the equilibrium path. Suppose now that seller i sells at $t = k$ along the equilibrium path. We show that the price that this seller sets, on the equilibrium path, must also equal 1. Let \hat{V}_C^{k-1} denote the buyer's equilibrium continuation value at the beginning of period $k - 1$. Since $\hat{V}_C^{k-1} = 0$, this implies that the any seller prices at 1 if the buyer deviates by not buying from this seller at any date $t \in \{1, 2, \dots, k - 1\}$. Consider now seller i who sells at date k on the

equilibrium path. If this seller also sells at every subsequent date, then he is a monopolist for all remaining periods on path, and will therefore price at 1. If not, let \hat{t} be the last period where seller $j \neq i$ makes a sale, and let this seller have m_j units of endowment at this point. By pricing at $1 - \varepsilon$ in periods $t \in \{1, 2, \dots, \hat{t}\}$, seller i makes a sale in every period, and can earn $m_j(1 - \varepsilon)$ for any $\varepsilon > 0$. Thus his equilibrium continuation value at the beginning of period k must equal \hat{V}_C^k , implying that he must price at 1 in period k .

Our second step is to show that in any pure strategy equilibrium, $d^t \in \{A, B\} \forall t$ along the equilibrium path, thereby completing the proof.

Suppose to the contrary that there is some period where $d^t = \emptyset$. Let k be the period with the smallest index such that $d^k = \emptyset$ on the equilibrium path. Let seller i have positive inventory at date 1. If both sellers have positive inventory, then the price on the equilibrium path at date 1 must equal zero for both sellers. Now if i offers a price $p_i^k \in (0, 1)$, then both i and C has a positive current payoff if C buys from i . Furthermore, j does not have an observable deviation, and thus j 's continuation strategy does not change. Thus $d^k = \emptyset$ cannot be optimal when both sellers have positive inventory at date 1.

If only seller i has positive inventory at date 1, then i has at least 2 units, and can therefore sell 1 unit at date k without reducing his continuation value. Thus $d^k = \emptyset$ cannot be optimal in this case as well. ■

Proof of lemma 8: 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. ■

References

- [1] Anton, J. J., & Yao, D. A. (1989). "Split awards, procurement, and innovation". *RAND Journal of Economics*, 538-552.
- [2] Anton, J. J., & Yao, D. A. (1992). "Coordination in split award auctions". *Quarterly Journal of Economics*, 107(2), 681-707.
- [3] Anton, J.J, G. Biglaiser, and N. Vettas. (2012). "Dynamic price competition with capacity constraints and a strategic buyer." Economic Research Initiatives at Duke (ERID) Working Paper 137.
- [4] Bagwell, K. (1995). "Commitment and Observability in Games", *Games and Economic Behavior*, **8**, 271-280
- [5] Bergemann, D., and J. Välimäki (1996). "Learning and Strategic Pricing", *Econometrica* **64**, 1125-1149.
- [6] Bergemann, D., and J. Välimäki (2003). "Strategic Buyers and Privately Observed Prices", *Journal of Economic Theory*.
- [7] Bergemann, D., and J. Välimäki (2006). "Dynamic Price Competition", *Journal of Economic Theory* 127, 232-263.
- [8] Bhaskar, V. (2009). "Commitment and Observability in a Contracting Environment", *Games and Economic Behavior* 66, 708–720.
- [9] Cabral, L. M., and M.H. Riordan. (1994). "The learning curve, market dominance, and predatory pricing". *Econometrica* 1115-1140.
- [10] Clarke, R., Davies, S., Dobson, P.W. and Waterson, M., (2002). *Buyer Power and Competition in European Food Retailing*, Edward Elgar, Cheltenham.

- [11] Compte, O. (2000). "Secret Price Cutting and Strategic Buyers", mimeo
- [12] Cooper, D.J., and R. G. Donaldson. (1998). "A Strategic Analysis of Corners and Squeezes", *Journal of Financial and Quantitative Analysis* **33**, 117-137.
- [13] van Damme, E., and S. Hurkens, (1997). "Games with Imperfectly Observable Commitment", *Games and Economic Behavior* **21**, 282-308.
- [14] Dobson P and M.Watson, (1999). "Retailer Power: Recent Developments and Policy Implications, *Economic Policy* 28, 135-
- [15] Dudey, M. (1992). "Dynamic Edgeworth-Bertrand Competition", *Quarterly Journal of Economics* **6**, 1461-1477.
- [16] Felli, L, and C. Harris, (1996). "Job Matching, Learning and the Distribution of Surplus", *Journal of Political Economy*.
- [17] Galbraith, J.K., (1952). *American Capitalism: The Concept of Countervailing Power*, Hamish Hamilton.
- [18] Ghemawat, P., and A. M. McGahan. (1998). "Order Backlogs and Strategic Pricing: The Case of the US Large Turbine Generator Industry", *Strategic Management Journal* **19**, 255-268.
- [19] Güth, W., G. Kirchsteiger and K. Ritzberger, (1998). "Imperfectly Observable Commitments in n-Player Games", *Games and Economic Behavior*, **23**, 54-74.
- [20] Inderst, R. (2008). "Single sourcing versus multiple sourcing". *RAND Journal of Economics*, 39(1), 199-213.
- [21] Inderst, R., and C. Wey, (2007). "Buyer Power and Supplier Incentives", *European Economic Review* 51, 647-667.

- [22] Lewis, T. R., and H. Yildirim. (2002). "Managing dynamic competition". *American Economic Review*, 92, 779-797.
- [23] Maggi, G.,(1999). "The Value of Commitment with Imperfect Observability and Private Information", *Rand Journal of Economics* **30**, 555-574.
- [24] Morgan, J., and F. Várdy, (2007). "The Value of Commitment in Contests and Tournaments when Observation is Costly", *Games and Economic Behavior* **60**, 326-338.
- [25] Rotemberg, J., and G. Saloner. (1986). "A Supergame Theoretic Model of Business Cycles and Price Wars during Booms", *American Economic Review* 76, 390-407.
- [26] Snyder, C. (1996) "A Dynamic Theory of Countervailing Power", *RAND Journal of Economics* **27**, 747-769.
- [27] Snyder, C. (1998). "Why do Larger Buyers Pay Lower Prices? Intense Seller Competition", *Economics Letters* 58, 205-209.
- [28] Stole, L., and J. Zwiebel. (1996). "Organizational Design and Technology Choice under Intrafirm Bargaining, *American Economic Review* 86, 195-222.
- [29] von Ungern-Sternberg, T. (1996). "Countervailing Power Revisited", *International Journal of Industrial Organization* 14, 507-519.