Center for Research on Economic and Social Theory

CREST Working Paper

Intertemporal Self-Selection with Multiple Buyers Under Complete and Incomplete Information

Mark Bagnoli
Stephen W. Salant
Joseph E. Swierzbinski

April, 1992
Number 92-05

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109-1220

Intertemporal Self-Selection with Multiple Buyers Under Complete and Incomplete Information[1]

Mark Bagnoli
Indiana University

Stephen W. Salant
University of Michigan

Joseph E. Swierzbinski
University of Michigan

April, 1992

[1]Please address all correspondence to Stephen W. Salant, Department of Economics, University of Michigan, Ann Arbor, Michigan 48109.

Abstract

We consider a monopolist selling durable goods to consumers with unit demands but different preferences for quality. The seller can offer items of different quality at the same time to induce buyers to self-select, as in Mussa-Rosen (1978), but is not artificially constrained to offer only one such menu. Instead the seller can offer without precommitment a *sequence* of menus over time. In the two-buyer case where the seller has complete information about each buyer's marginal valuation for quality, the seller's profits exceed what can be obtained from a single menu and sometimes approximate the profits of a perfectly discriminating monopolist. This result is no mere artifact of the assumption of complete information. As we show, even if the seller has incomplete information about the realized distribution of buyer types, he still may obtain higher expected profits (sometimes the entire surplus) by making a sequence of offers rather than a single offer.

1. Introduction

Consider the seller of a durable good. Suppose buyers want at most one item but differ in their willingness to pay for quality. Since buyers differ, the seller has an incentive to price discriminate; but assume he cannot offer an item to any buyer without making it available to everyone else. To price discriminate, the seller must induce buyers of one type to purchase a particular item while buyers of a different type purchase a different item.

In principle, the seller can induce such self-selection in two ways. At any given time, he can offer items of different quality which will have differing appeal to the heterogeneous buyers. Alternatively, the seller may be able to induce *intertemporal* self-selection since anticipation of his future offers may lead some types of buyers to purchase a current item while others prefer to wait for something more appealing.

Each of these forms of self-selection has been studied extensively in the literature—but mainly in isolation. For example, Coase's conjecture about the unprofitability of durable-goods monopoly spawned a literature on intertemporal self-selection; but since the monopolist in this literature offers items of only a single quality, he cannot induce self-selection at a point in time.[1] Mussa and Rosen (1978), on the other hand, examine self-selection at a point in time; but by constraining their monopolist to offer only a single menu of items, they eliminate the possibility of intertemporal self-selection.

We consider here the problem of a durable-goods monopolist who can make a *sequence* of offers to two heterogeneous buyers with unit demands but who cannot precommit. Since each buyer purchases at most one unit, the monopolist sells at most two units regardless of the number of offers he makes. We examine the case where the seller observes the valuation of each buyer and also the case where the seller is uncertain about the realized distribution of buyer types.

Suppose initially that the seller has complete information about the valuations of the two buyers and can make only a finite number of offers. If both buyers remain in the final period of this finite-horizon game, the equilibrium is identical to what Mussa-Rosen describe: in the absence of a corner solution, the monopolist finds it optimal to offer a menu of items to induce self-selection. The optimal menu necessarily contains items of inefficient quality and fails to extract the entire surplus.

In *every* period prior to the last, however, the monopolist finds it optimal to abandon menus. Instead, he offers a single item of *efficient* quality based on his observation of the marginal valuations of the remaining buyers.

Inducing self-selection over time turns out in this model to dominate inducing self-selection at a point in time. The monopolist offers a single item of a particular quality for sale and replaces it with an item of a different quality if and only if a sale occurs. Although the next item offered is more attractive to the previous purchaser than the item he actually purchased, that buyer is nonetheless rational to purchase the previous item. For he anticipates that the seller would not have made the subsequent item available until the previous item was purchased. By utilizing intertemporal self-selection—instead of point-in-time self-selection—the monopolist avoids having to distort the quality of the second item in order to make it *unattractive* to the first buyer; instead, he makes the second item *unavailable* (to either buyer) until the prior item is purchased.

The benefit of using intertemporal self-selection is that the monopolist can dispense with sales of inefficient quality. The cost is that he must space out his sales rather than make them at the same time. However, this cost becomes arbitrarily small for sufficiently high discount factors (i.e. a sufficiently short interval between offers).

For such discount factors, the monopolist can increase the present value of his profits relative to Mussa-Rosen's single-offer case and sometimes receives approximately the present value obtained by a perfectly discriminating monopolist.

These results differ so markedly from the received wisdom on principal-agent problems (whether interpreted as applying to second-degree price discrimination, optimal regulation, or optimal income taxation) that readers may wonder whether they are artifacts of some assumption about the time horizon, the number of buyers, or the information structure. For this reason, we also consider a two-buyer game with complete information where the monopolist can make an unlimited number of offers in succession. The equilibrium we consider has the same characteristics as in the finite-horizon case: no multi-item offers, no quality distortions, and, for discount factors sufficiently close to one, a present value for the monopolist that approximates the profits from perfect price discrimination. An equilibrium with these same characteristics exists for the infinite-horizon, complete-information game with *any* finite number of buyers.[2]

Other readers may wonder whether our results are artifacts of the assumption of complete information. For this reason, we show that our results persist in two-period, two-buyer games where the seller is uncertain about the realized distribution of buyer types. That a seller with incomplete information can benefit from making a sequence of offers instead of a single offer (and can sometimes achieve the expected profits of first-degree price discrimination) even for a discount factor of one may be surprising to readers familiar with the literature on noncooperative bargaining games. For in one-buyer games where the seller makes each offer and cannot observe the buyer's valuation, a common conclusion is that there is no advantage to making a sequence of offers instead of a single offer if the players' discount factor is sufficiently high.[3] The presence of more than one buyer sometimes creates opportunities for the seller which have no counterpart when he faces only one buyer. In the class of equilibria that we consider, the opportunity to make a sequence of offers increases the monopolist's profits only if there is some correlation between the buyers' types.

A clearer perspective on our results can be gained by viewing them in the context of a quite different literature. Arrow (1986) and Holmstrom (1982) have shown that a seller with complete information can extract the full surplus from buyers even when the seller is legally barred from making specific qualities unavailable to some buyers, provided he can condition the price paid for each item on the *collective* behavior of all buyers. Demski and Sappington (1984) and Cremer and McLean (1985) (among others) have shown that even with incomplete information, a seller can often extract the full surplus from risk-neutral buyers if prices are allowed to depend on the buyers' collective behavior and the buyers' types are correlated.

Arrow (1986) has criticized such contracts as involving "collective punishment." In contrast the contracts offered by the Mussa-Rosen monopolist do not involve collective punishment since the net utility which a buyer can achieve in a period does not depend on the *contemporaneous* behavior of the other buyers. If a seller makes a sequence of such offers, however, he will take prior buyer behavior into account when formulating his current offer. The payoff to one buyer will then depend on the *prior* behavior of the other buyer even if it does not depend on his contemporaneous behavior.

Our paper investigates under both complete and incomplete information the degree to which a sequence of offers—none of which involves contemporaneous collective punishment—serves as a substitute for a single offer involving such punishment. We identify cases under both complete and incomplete information where the monopolist extracts the full surplus; more generally, we find that the monopolist achieves a higher surplus using a sequence of offers than he does with a single offer.

In the next section, we examine the multi-offer, complete-information game with a finite or unbounded horizon. Section 3 extends our results to incomplete information and section 4 contains some concluding remarks.

2. Optimal Sequential Offers Under Complete Information

In this section, we consider the gains which the monopolist can achieve by making a *sequence* of offers to the two buyers instead of a single offer. Making a sequence of offers is advantageous because it induces buyers of different types to purchase in different periods. For discount factors sufficiently close to one, we will show that the monopolist can always increase the present value of his profits by making a succession of single-item offers.

We model the problem as a multi-period game, with the periods indexed by $t = 1, 2 \ldots T$. There are two stages in each period. In the first stage, the monopolist makes an offer consisting of a menu of items, with each item fully characterized by its quality level and price. In the second stage, each buyer simultaneously chooses whether to select a single item from among those offered or to continue to the next period. Buyers are assumed to have unit demands and select at most one item during the course of the game. Since the monopolist makes an offer in each period, the horizon, T, corresponds to the maximum number of distinct offers that the monopolist can make during the game. When T is infinite, the number of offers that the monopolist can make is unbounded.

If buyer i purchases an item of quality q_i in period t_i and pays price p_i, then he obtains the utility (discounted to $t = 1$):

$$U_i = \beta^{t_i - 1} (v_i q_i - p_i) \ , \tag{1}$$

where β is a discount factor assumed to be the same for all players and v_i is a positive constant. We assume that $1 > \beta > 0$ and that $v_1 > v_2 > 0$. A buyer who chooses not to purchase in any period receives a utility of zero.

If buyer 1 accepts (p_1, q_1) in period t_1 and buyer 2 accepts (p_2, q_2) in t_2, then the present value received by the monopolist is:

$$V = \beta^{t_1 - 1} (p_1 - c(q_1)) + \beta^{t_2 - 1} (p_2 - c(q_2)) \ . \tag{2}$$

If only the single buyer i purchases during the game (in period t_i), then the monopolist receives $\beta^{t_i - 1} (p_i - c(q_i))$, and if neither buyer makes a purchase then the monopolist receives zero. Assume that $c'(q)$, $c''(q) > 0$ for all q, that $c(0) = 0$, and that $v_2 > c'(0)$ [4].

The marginal valuations for quality, v_1 and v_2, are assumed to be common knowledge. Throughout this section, we also assume that the monopolist has complete information and knows the type (marginal valuation) of each specific buyer but that he is legally barred from using this information to make separate offers to each type of buyer.

A player's strategy in each period specifies his action in that period as a function of time and anything else that he can observe — for example, the sequence of previous offers, the number of buyers purchasing specific price-quality pairs, the periods in which various items were purchased, the marginal valuations of the buyers who remain, and, for the buyers, the menu of prices and qualities offered in the first stage of the current period. The monopolist's strategy in each period specifies his offer in the first stage of the period as a function of these things. Similarly, a buyer's strategy specifies which price-quality pair (if any) to accept in the second stage of each period. The monopolist chooses his strategy to maximize the present value of his profits. Each buyer chooses a strategy to maximize his own utility.

Mussa-Rosen (1978) assumed that the monopolist's first offer was also his last offer ($T = 1$). For purposes of comparison, we begin by briefly summarizing their findings. In the subgame-perfect equilibrium of this one-offer game, the monopolist adopts the following strategy. The monopolist offers (p_1, q_1) and (p_2, q_2) in order to:

$$\max \ (p_1 - c(q_1)) + (p_2 - c(q_2)) \tag{3a}$$

subject to the constraints:

$$\begin{aligned} p_i, q_i &\geq 0 \text{ for } i = 1, 2 \\ v_1 q_1 - p_1 &\geq 0 && \text{(3b)} \\ v_2 q_2 - p_2 &\geq 0 && \text{(3c)} \\ v_1 q_1 - p_1 &\geq v_1 q_2 - p_2 && \text{(3d)} \\ v_2 q_2 - p_2 &\geq v_2 q_1 - p_1 \ . && \text{(3e)} \end{aligned}$$

The optimal strategy for each buyer in the second stage is to choose the best item offered in the first stage that generates nonnegative utility. Equations (3b) and (3c) are "participation" constraints indicating that buyers will only purchase items which generate nonnegative utility. Equations (3d) and (3e) are "self-selection" constraints. Equation (3d) states that (p_1, q_1), the item intended for buyer 1, will only be purchased by buyer 1 if it generates utility for him that is at least as high as the utility generated by the alternative item. Equation (3e) has a similar interpretation for buyer 2.

The optimization problem specified by equations (3a) through (3e) has been extensively analyzed in the literature: the optimization has a unique solution; moreover, at the optimum, the self-selection constraint for buyer 1 and the participation constraint for buyer 2 are binding.

By solving the binding constraints (equations (3c) and (3d)) for p_1 and p_2 and substituting the resulting expressions in equation (3a), we obtain a strictly concave objective function in q_1 and q_2 which is maximized subject only to the nonnegativity constraints for q_1 and q_2. Solving the resulting optimization problem, we obtain for q_1: $v_1 - c'(q_1) = 0$. As for q_2, there are two possibilities: Either $q_2 > 0$ and $2v_2 - v_1 - c'(q_2) = 0$ or, alternatively, $q_2 = 0$ and $2v_2 - v_1 - c'(0) \leq 0$. In this "corner case," the interpretation of $q_2 = 0$ is that buyer 2 is not served. In either the interior or the corner case, the binding constraints (3c) and (3d) can be solved to yield the optimal prices (p_i).

If $2v_2 - v_1 - c'(0) > 0$, then it is optimal for the monopolist to sell to both buyers. Let (p_i^{MR}, q_i^{MR}) denote the item purchased by a type i buyer in this case. It will be convenient to define the following additional variables for the interior case: π_i^{MR}, V^{MR}, and U_i^{MR}. Let π_i^{MR} denote the profit the monopolist derives from the type i buyer when both types buy:

$$\pi_i^{MR} = p_i^{MR} - c(q_i^{MR}) . \tag{4}$$

Let V^{MR} denote the sum of the payoffs which the monopolist collects from the buyers in this case:

$$V^{MR} = \pi_1^{MR} + \pi_2^{MR} . \tag{5}$$

Let U_i^{MR} denote the equilibrium payoff of buyer i in this case. That is, $U_2^{MR} = v_2 q_2^{MR} - p_2^{MR} = 0$ and $U_1^{MR} = v_1 q_1^{MR} - p_1^{MR}$.

In contrast, recall that a *perfectly discriminating* monopolist with complete information is less heavily constrained. He can offer an item to one buyer while making it unavailable to the other buyer. Therefore his objective function is the same as (3a) but he is restricted *only* by the two participation constraints (3b and 3c)—and not by either self-selection constraint. The profit obtained from each type i buyer will be $v_i q_i - c(q_i)$. The optimal item for such a monopolist to offer to a type i buyer has a price and quality level denoted by (p_i^{1st}, q_i^{1st}) and specified by the equations:

$$p_i^{1st} = v_i q_i^{1st} , \tag{6a}$$
$$c'(q_i^{1st}) = v_i . \tag{6b}$$

Since the monopolist can extract all of the surplus from each buyer, he offers each the "efficient" quality level (which maximizes total surplus).

Let π_i^{1st} denote the profit that the monopolist obtains when a type i buyer accepts the optimal take-it-or-leave-it offer specified in equations (6a) and (6b).

$$\pi_i^{1st} = p_i^{1st} - c(q_i^{1st}) . \tag{7}$$

The buyer with the higher marginal valuation for quality (buyer 1) always receives the same quality level from the Mussa-Rosen monopolist as he would from a perfectly discriminating monopolist. Since the offer of the perfectly discriminating monopolist would violate the self-selection constraint of buyer 1, the Mussa-Rosen monopolist can never extract the entire surplus.

We now consider the case where $T > 1$. We consider in turn the two parameter regimes which gave rise in Mussa-Rosen's single-offer case, respectively, to corner and interior equilibria. In each case, we begin by describing the subgame-perfect equilibrium strategies for each of the three players. Having described them, we then verify that each strategy is optimal in every subgame given the other strategies.

In the "corner" regime where $2v_2 - v_1 - c'(0) \leq 0$, the following strategy combination forms a subgame-perfect equilibrium in both the finite- and the infinite-horizon games. In each period, each buyer accepts the best offer (if any) of those which generate nonnegative utility for him. We refer to this strategy as the "get-it-while-you-can" strategy since the buyer seizes the first opportunity for surplus as if no future opportunities might present themselves.

The monopolist uses the following strategy in period t. If buyer 1 or both buyers remain in the game at the beginning of the period, then the monopolist offers the single price-quality pair (p_1^{1st}, q_1^{1st}) defined by equations (6a) and (6b). If only buyer 2 remains, then the monopolist offers the single price-quality pair (p_2^{1st}, q_2^{1st}).

Next, we verify that these strategies are subgame perfect. In both finite- and infinite-horizon games, no buyer can ever obtain positive utility when the monopolist uses the proposed strategy. Hence, it is optimal for each buyer to use the get-it-while-you-can strategy in each period.

Finally, we verify the optimality of the seller's strategy when buyers use the get-it-while-you-can strategy in each period. In any period with only one remaining buyer, the proposed strategy is clearly optimal for the monopolist since it extracts the maximum possible surplus from this buyer.

Suppose instead that both buyers remain. Any price-quality pair (with $q > 0$) that generates nonnegative utility for buyer 2 generates strictly positive utility for buyer 1. Hence, it is infeasible—given the buyers' strategies—for the monopolist to sell only to buyer 2. This leaves the monopolist two choices: sell immediately to both buyers or sell only to buyer 1. Among all the offers which only buyer 1 would accept, the proposed offer is clearly optimal for the monopolist since it extracts the maximum possible surplus from buyer 1. Moreover, since for this parameter regime a corner solution is optimal in the single-period game, the profit from selling only to buyer 1 must strictly exceed the profit from selling simultaneously to the two buyers.

Indeed, if at least one future period remains the proposed strategy is even more profitable than in the single-offer case since the opportunity to make future offers enables the seller to extract the entire surplus of buyer 2 in the next period. Hence, the proposed sequential strategy results in the following present value for the monopolist:

$$V^S = \pi_1^{1st} + \beta\pi_2^{1st} . \tag{8}$$

The above arguments demonstrate that when $2v_2 - v_1 - c'(0) \leq 0$, the proposed strategy combination constitutes a subgame-perfect equilibrium for the complete-information game with any horizon length, either finite or infinite. Moreover, these strategies constitute a subgame-perfect equilibrium for any discount factor between zero and one. In any game in which the monopolist can make two or more offers ($T \geq 2$), the monopolist receives the equilibrium present value specified in equation (8). This present value strictly dominates the profit obtained by a static Mussa-Rosen strategy and approaches $\pi_1^{1st} + \pi_2^{1st}$, the profit obtained by a perfectly discriminating monopolist, as the discount factor approaches one.

We now consider the "interior" regime where $2v_2 - v_1 - c'(0) > 0$. Since we are ultimately interested in the play of the game when the time interval between successive offers becomes negligible, we focus on the equilibrium strategies when the discount factor pertaining to the interval between offers approaches one.

For discount factors sufficiently close to one, the subgame-perfect equilibrium strategies in the finite-horizon game may be described as follows. Buyer 2 uses the get-it-while-you-can strategy in each period. Buyer 1 uses the get-it-while-you-can strategy in any subgame where he is the sole remaining buyer. In addition, buyer 1 always uses the get-it-while-you-can strategy in the final period of the game ($t = T$) whether or not buyer 2 remains. Finally, if both buyers remain prior to T, buyer 1 uses the get-it-while-you-can strategy in the current period whenever he anticipates being the only remaining buyer from the next period onward; buyer 1 will anticipate being isolated if he observes that some item currently offered by the monopolist will provide nonnegative utility for buyer 2.

If both buyers remain prior to T and if buyer 1 observes that no item offered will yield nonnegative utility to buyer 2, then buyer 1 is more selective: in the current period t he will accept the best price-quality pair (if any) of those that generate utility (discounted to period 1) greater than or equal to the reservation utility (also discounted to period 1) $\beta^{T-1}U_1^{MR}$. Buyer 1's strategy is a generalization of the get-it-while-you-can strategy for the case where the discounted reservation utility is $\beta^{T-1}U_1^{MR}$.

As for the description of the seller's strategy, he will offer a single item of efficient quality priced so as to extract the entire surplus in any period where only one buyer remains. If both buyers remain in the final period, the monopolist will offer the Mussa-Rosen menu. If both buyers remain prior to T, the monopolist will offer the single quality q_1^{1st} at a price $p(t)$ which is given by the equation:

$$p(t) = (1 - \beta^{T-t})p_1^{1st} + \beta^{T-t}p_1^{MR} . \tag{9}$$

In equilibrium, buyer 1 accepts $(p(1),\ q_1^{1st})$ in the first period and buyer 2 accepts $(p_2^{1st},\ q_2^{1st})$ in the second period.

To verify that these strategies are subgame perfect, note first that the get-it-while-you-can strategy is optimal for buyer 2 since he can never obtain positive utility when the monopolist uses his equilibrium strategy. For the same reason, the get-it-while-you-can strategy is optimal for buyer 1 when he is the sole remaining buyer. Since a failure to accept an offer in the final period results in zero utility, the get-it-while-you-can strategy is also optimal for buyer 1 in the final period whether or not buyer 2 remains. Finally, if the monopolist offers a price-quality pair that provides nonnegative utility for buyer 2, then buyer 2, using the get-it-while-you-can strategy, will accept it; buyer 1 can therefore expect only zero utility in future periods, and the get-it-while-you-can strategy is also optimal for him in this case.

If both buyers remain prior to T and the monopolist offers no item in period t which gives buyer 2 nonnegative utility, then buyer 1 anticipates that buyer 2 will remain and that the seller will offer q_1^{1st} in the next period at a price of $p(t+1)$ defined in equation (9). For all $t+1 < T$, equation (9) implies that $p(t+1) > p_1^{MR}$. Hence, the previous analysis of the single-offer game indicates that buyer 2 will reject such an offer. More generally, buyer 1 can anticipate that as long as he declines to purchase, the monopolist will continue to offer $(p(s),\ q_1^{1st})$ in any period $t < s < T$ and that buyer 2 will continue to reject these offers. Moreover, by substituting equations (6a) and (9) into equation (1) and using the fact that $U_1^{MR} = p_1^{1st} - p_1^{MR}$, it is straightforward to verify that buyer 1 obtains the same discounted utility, $\beta^{T-1}U_1^{MR}$, by accepting any of the offers $(p(s), q_1^{1st})$.

Finally, suppose that buyer 1 declines to purchase until the final period. In this case, he can anticipate that as his final act the monopolist would offer the Mussa-Rosen menu and that buyer 1 would again earn the discounted utility $\beta^{T-1}U_1^{MR}$ by accepting an item from this menu. Hence, in any future period including the last, the buyer expects to obtain the discounted utility $\beta^{T-1}U_1^{MR}$ if he fails to accept the monopolist's offer in period t when both buyers remain. It is thus optimal for buyer 1 to accept the best offer in period t of those that provide discounted utility greater than or equal to $\beta^{T-1}U_1^{MR}$; that is, it is optimal for buyer 1 to use the generalized get-it-while-you-can-strategy with discounted reservation utility $\beta^{T-1}U_1^{MR}$.

To complete the verification of subgame perfection, we consider the optimality of the seller's strategy. As in the other parameter regime, the proposed strategy is clearly optimal for the monopolist in any subgame with one remaining buyer since it extracts the maximum possible surplus from such a buyer. Moreover, as previously noted, the Mussa-Rosen menu is optimal for the monopolist in the final period when both buyers remain.

If two buyers remain prior to the final period, the monopolist has two choices just as in the previous parameter regime: sell immediately to both buyers or sell only to buyer 1. We now verify that, for sufficiently high discount factors, the seller will *never* make an offer containing two items prior to the final period. To sell simultaneously to both buyers, the monopolist must provide buyer 2 with nonnegative utility. But this will make buyer 1 anticipate isolation and he too will accept the best offer generating nonnegative utility. The analysis of the single-offer game demonstrates that when both buyers require nonnegative utility to participate and $2v_2 - v_1 - c'(0) > 0$, the static Mussa-Rosen offer is the most profitable two-item offer for the seller and results in a payoff of $\pi_1^{MR} + \pi_2^{MR}$.

The monopolist can easily dominate this best two-item offer provided he can make at least one future offer (that is, if $t < T$). For, suppose the monopolist simply deleted the lower quality item from the Mussa-Rosen offer. Buyer 2 will not accept the remaining item since it provides him with negative utility (as it did when it was part of the Mussa-Rosen menu). On the other hand, buyer 1 can obtain the discounted utility of $\beta^{t-1}U_1^{MR}$ by accepting the remaining item in period t. Since $\beta^{t-1}U_1^{MR} > \beta^{T-1}U_1^{MR}$, buyer 1—using his equilibrium strategy—would purchase the item. Hence, the monopolist can anticipate that buyer 1 will accept the offer and buyer 2 will reject it. The seller could then offer (p_2^{1st}, q_2^{1st}) next period and, since it would provide buyer 2 with nonnegative utility, could anticipate that the offer would be accepted. This strategy yields the seller $\pi_1^{MR} + \beta\pi_2^{1st}$. For $\beta > \beta^*$ where $\beta^* = \pi_2^{MR}/\pi_2^{1st}$ this payoff strictly exceeds the payoff from the two-item Mussa-Rosen offer.

Of course, this strategy—while superior to *any* two-item menu—is not the best single-item offer. Recall that in any period t prior to the last period ($t < T$), if both buyers remain and buyer 1 does not anticipate that buyer 2 will make a purchase in period t, then buyer 1 will accept the best offer yielding discounted utility $\beta^{T-1}U_1^{MR}$. The monopolist should therefore continue to offer q_1^{1st} since that maximizes the total surplus; but he should price it so as to extract all but the minimum surplus necessary for buyer 1 to purchase. The price of the proposed dominating offer (p_1^{MR}) is suboptimal since it gives U_1^{MR} to buyer 1 when he requires, in terms of period t utility, only $\beta^{T-t}U_1^{MR}$. The optimal price should therefore be $p(t) = p_1^{MR} + (U_1^{MR} - \beta^{T-t}U_1^{MR})$. Using the fact that $U_1^{MR} = p_1^{1st} - p_1^{MR}$, we conclude that the optimal price for the item of quality q_1^{1st} is given by equation (9).

The above discussion demonstrates that the proposed strategies constitute a subgame-perfect equilibrium in the two-buyer, finite-horizon game with complete information for any $\beta > \beta^*$. In equilibrium, the monopolist sells a unit of quality q_1^{1st} to buyer 1 in the first period at a price given by $p(1)$ in equation (9) and makes the offer (p_2^{1st}, q_2^{1st}) to buyer 2 in the second period. Substituting these prices and qualities into equation (2) and collecting terms using equations (4) and (7), we obtain the following expression for the equilibrium present value obtained by the monopolist in the play of the T-period game:

$$V_T = (1 - \beta^{T-1})\pi_1^{1st} + \beta^{T-1}\pi_1^{MR} + \beta\pi_2^{1st} . \tag{10}$$

For infinite-horizon games with complete information and sufficiently high β, the analysis when $2v_2 - v_1 - c'(0) > 0$ is virtually identical to the analysis for the previous parameter regime. In the infinite-horizon game, the use in each period of the get-it-while-you-can strategy by each buyer and the strategy where the monopolist offers a single-item menu which extracts the entire surplus of the remaining buyer with the highest marginal valuation constitutes a subgame-perfect equilibrium. The equilibrium present value obtained by the monopolist, V^S, is given by equation (8).

As in the previous parameter regime, the get-it-while-you-can strategy is optimal for each buyer since the buyers can anticipate only zero utility in future periods. The strategy proposed for the seller is optimal if $V^S > V^{MR}$. A comparison of equations (5) and (8) indicates that $V^S > V^{MR}$ if $\beta > \beta^*$. In the play of the equilibrium, the monopolist makes the offer (p_1^{1st}, q_1^{1st}) in the first period.

It can be verified by working backwards that whenever the horizon is finite the subgame-perfect equilibrium is essentially unique.[5] In contrast, there are a continuum of equilibria in the infinite-horizon case. The particular infinite-horizon equilibrium we discussed was chosen because the strategies supporting it are the limiting strategies in the finite-horizon case for a sufficiently long horizon. Buyer 2 in either case accepts the best offer in a period yielding nonnegative utility. Buyer 1 in the infinite-horizon case also accepts such offers and in the finite-horizon case adopts the identical strategy prior to T except when two buyers remain and no current item would give buyer 2 nonnegative utility. In such circumstances, buyer 1 purchases the best item offered provided it yields $\beta^{T-1}U_1^{MR}$. But this strategy becomes virtually the same as "get-it-while-you-can" as T becomes large. Recall that when the monopolist is restricted to a maximum of T offers, he makes the offer $(p(1), q_1^{1st})$ in the first period, with $p(1)$ given in equation (9), and obtains the equilibrium present value, V_T, given in equation (10). For each fixed value of β greater than β^*, equations (8), (9), and (10) indicate that the monopolist's first-period offer and present value in the T-period game approach the infinite-horizon first-period offer, (p_1^{1st}, q_1^{1st}), and the present value obtained in the infinite-horizon game as T becomes large. ($p(1) \to p_1^{1st}$ and $V_T \to V^S$ as $T \to \infty$.)

As noted previously, V^S approaches the profit obtained by a perfectly discriminating monopolist as the discount factor approaches one.

3. Extension to Incomplete Information

We now consider a two-buyer game where the buyer types are random variables drawn from a distribution that is common knowledge. As in the previous section, we assume that there are only two possible buyer types (1 or 2) distinguished by their different marginal valuations for quality, $v_1 > v_2 > 0$. Let Θ_{11} be the prior probability of the event that both buyers are type 1, Θ_{22} the prior probability that both buyers are type 2, and Θ_{12} the probability that there is one buyer of each type. The seller does *not* know the realized distribution of buyer types. In order to focus on games with incomplete information, we assume that $\Theta_{ij} \neq 1$ for all i and j. For simplicity, we also restrict attention to one-period and two-period games.

As we will see, the introduction of incomplete information alters none of our main conclusions. Provided there is correlation between the buyer types, the monopolist's (expected) profit once again may increase when he is given the opportunity to offer a sequence of menus. Indeed, offering a sequence of menus permits the seller in some circumstances, to extract the entire (expected) social surplus.

Making offers in sequence permits the monopolist to use information about one buyer gleaned by observing the prior behavior of the other buyer. Such an opportunity cannot arise in the two special cases which currently inform the literature on sequential bargaining with one-sided information—the case with only one buyer and the case with multiple buyers but independent types.

In the perfect Bayesian equilibria we investigate, the monopolist offers a single item in the first period intended for type 1 buyers exactly as in section 2. He then offers a menu of items in the second period as he does in unreached subgames of the complete-information case.

As before, we begin with the single-offer game. In a game with either one or two buyers where the seller does not know the realized distribution of buyer types, the self-selection and participation constraints, equations (3b) through (3e), must continue to hold. For the menu item $(p_i,\, q_i)$ intended for a buyer of type i $(i = 1, 2)$ would be accepted by that buyer if and only if it gives him at least as much utility as he would get from purchasing the other item on the menu or purchasing nothing at all. As in section 2, these constraints imply the following relations between the prices and quantities in any profit-maximizing menu where q_i is intended for a type i buyer.

$$p_1 = v_1 q_1 - (v_1 - v_2) q_2 \;, \tag{11}$$

and

$$p_2 = v_2 q_2 \;. \tag{12}$$

Although the self-selection and participation constraints remain the same, the objective function in equation (3a) must now be replaced by the expected profit of the seller:

$$\bar{n}_1\big(p_1 - c(q_1)\big) + \bar{n}_2\big(p_2 - c(q_2)\big), \tag{13}$$

where $\bar{n}_i$ denotes the expected number of buyers of each type.[6]

Let q_i^{1st} and $p_i^{1st} = v_i q_i^{1st}$ denote the quality and price which a perfectly discriminating monopolist would offer to a type i buyer. As in section 2, it is always optimal for the Mussa-Rosen monopolist to offer a type 1 buyer the quality q_1^{1st} at a price specified by equation (11). Replacing $v_1 q_1$ by p_1^{1st} in equation (11) and substituting the resulting expression in equation (13) produces the following expression for the monopolist's objective function:

$$\bar{n}_1\Big(\pi_1^{1st} - [v_1 - v_2]q_2\Big) + \bar{n}_2\Big(v_2q_2 - c(q_2)\Big), \qquad (14)$$

where, as before, $\pi_i^{1st} = v_iq_i^{1st} - c(q_i^{1st})$ denotes the profits obtained from a type i buyer by a perfectly discriminating monopolist. Let W^{MR} denote the expected profits obtained by a Mussa-Rosen monopolist. An expression for W^{MR} is obtained by substituting the optimal value for q_2 in equation (14).

In the rest of this section we restrict attention to two-period, two-buyer games and perfect Bayesian equilibria of the following form.[7] In the first period, the seller makes an offer which would give type 2 buyers negative utility and hence does not interest them. Let u be the utility a type 1 buyer would get from the best item on such a menu. The type 1 buyer may accept this item always, reject this item always, or accept this item with some probability. He will randomize if and only if he expects a utility of u as well in the second period. Let α denote the probability that each type 1 buyer accepts the initial offer. In the equilibria we consider, the seller initially offers the single item with quality q_1^{1st} and price $p_1^{1st} - u$. In the second period, the seller offers a Mussa-Rosen menu based on his knowledge of α, his observation of whether or not someone purchased one item in the first period, and which item, if any, was purchased.[8]

Let $q_2(1,\alpha)$ and $q_2(2,\alpha)$ denote the qualities which the monopolist intends for a type 2 buyer in the second period when one or two buyers remain and type 1 buyers accept the first-period offer with probability α. Combining the probability of each sequence of purchases with the profits obtained in each case produces the following expression for the expected present value of the seller in the equilibrium under consideration:

$$\begin{aligned}
E\Pi \;=\; & +2\Theta_{11}\alpha^2\Big\{\pi_1^{1st} - u\Big\} \\
& +2\Theta_{11}(1-\alpha)^2\beta\Big\{\pi_1^{1st} - (v_1 - v_2)q_2(2,\alpha)\Big\} \\
& +2\Theta_{11}\alpha(1-\alpha)\Big(\Big\{\pi_1^{1st} - u)\Big\} + \beta\Big\{\pi_1^{1st} - (v_1 - v_2)q_2(1,\alpha)\Big\}\Big) \\
& +\;\Theta_{12}\alpha\Big\{\pi_1^{1st} - u + \beta\Big[v_2q_2(1,\alpha) - c\Big(q_2(1,\alpha)\Big)\Big]\Big\} \\
& +\Theta_{12}(1-\alpha)\beta\Big\{\pi_1^{1st} - (v_1 - v_2)q_2(2,\alpha) + v_2q_2(2,\alpha) - c\Big(q_2(2,\alpha)\Big)\Big\} \\
& +2\Theta_{22}\beta\Big\{v_2q_2(2,\alpha) - c\Big(q_2(2,\alpha)\Big)\Big\}. \qquad (15)
\end{aligned}$$

The quantity β is the discount factor and the quantity $\pi_1^{1st} - u$ is the profit obtained from each type 1 buyer who accepts a first-period offer. $E\Pi$ depends on both u and α.

Now consider equilibria where the type 1 buyers mix with strictly positive probability in the first period and so must be indifferent between accepting and rejecting the single item offered in the first period. If a type 1 buyer accepts, he gets u with certainty; if he rejects, he gets either $\beta(v_1 - v_2)q_2(1,\alpha)$ if the other buyer accepts the initial offer or $\beta(v_1 - v_2)q_2(2,\alpha)$ if the other buyer also declines it. (Equation (11) implies that a type 1 buyer receives $(v_1 - v_2)q_2$ from the Mussa-Rosen offer.) Since the type 1 buyer is indifferent between accepting and rejecting in the first period, it must be the case that:

$$u = \frac{2\alpha\Theta_{11}}{2\Theta_{11}+\Theta_{12}}\beta(v_1 - v_2)q_2(1,\alpha) + \left(1 - \frac{2\alpha\Theta_{11}}{2\Theta_{11}+\Theta_{12}}\right)\beta(v_1 - v_2)q_2(2,\alpha), \qquad (16)$$

where $2\alpha\Theta_{11}/(2\Theta_{11}+\Theta_{12})$ is the probability that a type 1 buyer assigns to the event that the "other" buyer makes a purchase.

We begin with a negative but illuminating result. We show for the special case where the seller would offer the *same* menu in the second period whether or not a purchase occurred in the first period that making a sequence of offers confers no advantage in the class of equilibria under consideration. Denote the quality intended for type 2 buyers in the second period regardless of whether one or two buyers remain as $\tilde{q}_2(\alpha)$.

If $\alpha = 1$, then $\tilde{q}_2(\alpha) = q_2^{1st}$ since (p_2^{1st}, q_2^{1st}) is the optimal second-period offer when only type 2 buyers remain. In this case, $u = \beta(v_1 - v_2)q_2^{1st}$ since this is the minimum utility that will induce a type 1 buyer to purchase in the first period. If $0 < \alpha < 1$, then equation (16) implies that $u = \beta(v_1 - v_2)\tilde{q}_2(\alpha)$. In either case, substituting $q_2(1,\alpha) = q_2(2,\alpha) = \tilde{q}_2(\alpha)$ and $u = \beta(v_1 - v_2)\tilde{q}_2(\alpha)$ in equation (15) and collecting the terms multiplying π_1^{1st}, $(v_2\tilde{q}_2(\alpha) - c(\tilde{q}_2(\alpha)))$, and $(v_1 - v_2)\tilde{q}_2(\alpha)$ produces the following expression for the expected profit of the monopolist:

$$\begin{aligned} E\Pi \;=\; & \beta\Big\{\bar{n}_1\big(\pi_1^{1st} - (v_1 - v_2)\tilde{q}_2(\alpha)\big) + \bar{n}_2\Big(v_2\tilde{q}_2(\alpha) - c\big(\tilde{q}_2(\alpha)\big)\Big)\Big\} \\ & + (1-\beta)\big[\bar{n}_1\alpha\pi_1^{1st}\big]. \end{aligned} \qquad (17)$$

Equation (14) implies that the sum in braces in equation (17) is the expected profit in a single-period game where the monopolist offers the most profitable feasible menu which provides quality q_1^{1st} to a type 1 buyer and $\tilde{q}_2(\alpha)$ to a type 2 buyer. The term in square brackets in equation (17) is less than or equal to $\bar{n}_1\pi_1^{1st}$ which is the expected profit that a monopolist in a single-period game can obtain by offering the single-item menu (p_1^{1st}, q_1^{1st}). Since the Mussa-Rosen offer is optimal in a one-period game, the expected profit which it generates, W^{MR}, is at least as large as that generated by any other feasible menu including either of the two menus mentioned above. Since $E\Pi$ is a beta-weighted average of two terms, neither of which strictly exceeds W^{MR}, the seller gains no advantage (for any admissible β) from the opportunity to make a second offer.

The Mussa-Rosen menu when either one or two buyers remain in the second period is determined by maximizing the objective function in equation (14) but with $\bar{n}_i$ updated in accordance with Bayes' rule. Hence, when the expected number of type 2 buyers is strictly positive, the optimal menu depends on the monopolist's beliefs *only* through the ratio of the (updated) expected number of type 1 buyers to the (updated) expected number of type 2 buyers.

Now consider a two-period game where the types of the two buyers are drawn independently from some (two-point) distribution and the monopolist offers a single item in the first period that type 1 buyers accept with probability $\alpha > 0$ (and type 2 buyers reject). Since we have excluded the degenerate case where $\Theta_{11} = 1$, the expected number of type 2 buyers in the second period will be strictly positive. Given independence, the ratio of the expected number of type 1 buyers to the expected number of type 2 buyers is the same whether one or two buyers remain in the second period.[9] Hence, the optimal menu is the same in these two cases. Together with the previous result, this shows that when the types of the two buyers are independent, the monopolist cannot increase his profits by engaging in the sort of intertemporal price discrimination which we are considering.

When it is optimal for the monopolist to condition the menu he offers in the second period on whether or not a purchase has occurred in the first period, intertemporal price discrimination may be advantageous relative to the point-in-time price discrimination of Mussa-Rosen. We now consider two examples where the monopolist does increase his profits by making a sequence of offers. In our first example, the buyers' types are perfectly, positively correlated. Our second example involves buyers with negatively correlated types where by a suitable choice of parameters any correlation coefficient between -1 and 0 can be considered.[10] For simplicity, we also assume in these examples that $\beta = 1$.

In our first example, $\Theta_{12} = 0$ but $\Theta_{11} > 0$ and $\Theta_{22} > 0$. The seller knows that both buyers are the same type but is uncertain which type that is. The equilibrium under consideration involves pure strategies and results in complete extraction of the expected social surplus.

We begin by summarizing the strategies and beliefs of the players. Each buyer believes the other buyer is the same type as himself. In each period, each type of buyer accepts the best available offer provided it yields nonnegative utility. In the first period, the monopolist offers (p_1^{1st}, q_1^{1st}). If no one purchases, the seller believes that he faces two type 2 buyers; if one buyer purchases, the seller believes that this buyer was type 1 and hence infers that the other buyer is also type 1. In the second period, if he observes 1 buyer remaining he repeats his initial offer; if he observes two buyers remaining, he offers (p_2^{1st}, q_2^{1st}).

We now turn to the optimality of this strategy combination. We consider first the optimality of the buyers' strategies. In general, if a buyer anticipates that he will receive zero utility in the future given the seller's strategy, then it is optimal for him to accept the best item currently offered provided it yields nonnegative utility. For this reason, the strategy of type 2 buyers is optimal in either period and the strategy of type 1 buyers is optimal in the second period.

Consider next the optimality of the strategy of a type 1 buyer in the first period. If the best item provides negative utility, then it is optimal for him to reject the offer. If the best item provides nonnegative utility, however, it is optimal that a type 1 buyer accept it. For, if he rejects it unilaterally, he anticipates that the other buyer—whom he infers is a type 1 like himself—will purchase in the first period and upon seeing this single purchase, the seller will offer him (p_1^{1st}, q_1^{1st}) in the next period. Since such an offer would generate zero utility, it is optimal for each type 1 buyer to accept the best item in the first period generating nonnegative utility.

We consider now the strategy and beliefs of the seller. If the seller offers (p_1^{1st}, q_1^{1st}) in the first period and no purchase occurs, the seller must conclude that both buyers are type 2. Given that belief, his strategy of offering (p_2^{1st}, q_2^{1st}) is clearly optimal.

If each buyer follows his equilibrium strategy, then the seller should never observe exactly one purchase of his first-period offer (p_1^{1st}, q_1^{1st}). If the seller does encounter this situation nonetheless, we postulate that the seller would conclude that a type 1 buyer purchased the initial item and that the remaining buyer is, therefore, also a type 1. Although a seller would not be able to account for the failure of a type 1 buyer to purchase initially as his strategy dictates, at least such an error by a type 1 buyer does not reduce the buyer's payoff. In contrast, if a type 2 buyer had purchased the initial item, his payoff would have been negative. Given the seller's belief that the remaining buyer is type 1, his strategy of repeating the offer of (p_1^{1st}, q_1^{1st}) after a single purchase is clearly optimal.

There is no need to verify that the initial offer of the monopolist is optimal since the proposed strategy achieves the expected profits of perfect price discrimination and no strategy could do better:

$$E\Pi = \Theta_{11} 2\pi_1^{1st} + \Theta_{22} 2\pi_2^{1st}.$$

For the same reason, there is no need to specify the beliefs of the seller or the strategies of players if the monopolist fails to offer (p_1^{1st}, q_1^{1st}) in the first period.

In our second example, we assume that $\Theta_{11} = 0$. Hence, the seller does not know whether he faces one or two buyers of type 2. However, the analysis below is easily extended to the case where $\Theta_{11} > 0$ and $\beta \leq 1$.

As in the case of complete information, there are "interior" and "corner" parameter regimes which determine whether it is optimal for the monopolist in a single-period game to serve both types of buyer or only type 1 buyers. As was also the case for the finite-horizon, complete-information game, the analysis of our second example is somewhat simpler for the corner case but the results are similar for the two parameter regimes. In the two-period, corner case, which is discussed below, it is straightforward to show that the monopolist can always obtain an expected profit greater than the static Mussa-Rosen profit by making a sequence of offers. In the two-period, interior case, the monopolist's optimal strategy in the first period involves maximizing the appropriate analog of the objective in equation (21) below. Numerical solutions of this optimization for the interior parameter regime confirm that there is a wide range of parameters for which the monopolist can increase his expected profit by making a sequence of offers.

In accord with the discussion in the previous paragraph, we assume for simplicity in what follows that:

$$1 > \Theta_{12} > V = \frac{2\,(v_2 - c'(0))}{v_1 - c'(0)}\,. \tag{18}$$

When this inequality holds, the optimal strategy for the monopolist in a single-period game is to serve only the type 1 buyer ($q_2 = 0$); hence, equation (14) implies that a Mussa-Rosen monopolist would earn the expected profit $W^{MR} = \Theta_{12}\pi_1^{1st}$.[11]

The following notation will be useful. Suppose that the monopolist makes a first-period offer which a type 2 buyer would reject and a type 1 buyer would accept with probability α. Let $\Theta_{12}(\alpha)$ denote the updated probability that there is one buyer of each type after both buyers are observed to reject this offer. From Bayes' rule it follows that:

$$\Theta_{12}(\alpha) = \frac{\Theta_{12} - \alpha\Theta_{12}}{1 - \alpha\Theta_{12}} . \tag{19}$$

Define α_0 via the equation:

$$\alpha_0 = \frac{1}{\Theta_{12}} \frac{\Theta_{12} - V}{1 - V} , \tag{20}$$

where V is the ratio defined in equation (18). α_0 is the largest value of α for which the inequalities in equation (18) are still (weakly) satisfied when Θ_{12} is replaced by $\Theta_{12}(\alpha)$. Intuitively, if type 1 buyers accept first-period offers with sufficiently small probability ($\alpha \leq \alpha_0$), then the observation that two buyers remain will still lead the seller in the second period to make a single-item offer intended only for the type 1 buyer. Equations (18) and (19) imply that α_0 is strictly greater than 0 and strictly less than 1.

Finally, we shall call a menu of items "type II" if it contains an item that provides nonnegative utility to type 2 buyers and "type I" if it does not contain such an item.

We begin by summarizing the equilibrium strategies of the two buyers. In each period, type 2 buyers purchase the best item offered if it yields nonnegative utility. A type 1 buyer adopts this same strategy in the second period and also in the first period when the menu offered is type II. Note that a type 2 buyer always accepts some item from a type II menu but never accepts an item from a type I menu.

Faced with a type I menu in the first period, a type 1 buyer's strategy depends on the maximum utility (denoted u) that he can obtain by accepting an item in the first period. If $u < 0$, then the type 1 buyer rejects the seller's initial offer. Recall that a type 1 buyer receives the utility $(v_1 - v_2)q_2$ from a Mussa-Rosen offer when q_2 is the quality of the item intended for type 2 buyers. Let $u_{max} = (v_1 - v_2)q_2^{1st}$. If $u > u_{max}$, then a type 1 buyer is assumed to accept the item that provides him with u. In the event that several items provide this maximum utility, assume that he chooses from among them the item of highest quality. For any maximum utility u, we refer to the highest quality menu item that provides a type 1 buyer with utility u as the "best item" for the type 1 buyer. If $u \in [0, u_{max}]$, then the type 1 buyer accepts the best item on a type I menu with probability $\alpha(u)$ (and rejects all other items), where $\alpha(u)$ is a continuous, strictly increasing function defined in the following paragraph .[12]

Let $q_2(2, \alpha)$ denote the quality that a Mussa-Rosen monopolist intends for type 2 buyers in a two-buyer game where there is at most one type 1 buyer and the probability that there is exactly one buyer of each type is given by $\Theta_{12}(\alpha)$. For $0 \leq u \leq u_{max}$, let $\alpha(u)$ denote the largest value of α that satisfies the equation: $u = (v_1 - v_2)q_2(2, \alpha)$. It is straightforward to verify that $q_2(2, \alpha)$ is a continuous function of α on the interval $[\alpha_0, 1]$ and that $\alpha(u)$ is a continuous, strictly increasing function of u on the interval $[0, u_{max}]$.[13] $\alpha(0) = \alpha_0$ and $\alpha(u_{max}) = 1$.

We turn now to the equilibrium strategy of the seller. The seller's strategy in the first period is defined in terms of the following construction. For $0 \leq u \leq u_{max}$, suppose that the seller were to offer the single item $(p_1^{1st} - u,\ q_1^{1st})$ in the first period and in the second period **i.** offer $(p_2^{1st},\ q_2^{1st})$ after observing a single purchase, and **ii.** after observing no purchases, offer the Mussa-Rosen menu which is optimal when the probability that there is one buyer of each type is given by $\Theta_{12}(\alpha(u))$. If the buyers follow their equilibrium strategies, then $E\Pi(u)$, the expected profit that would be obtained by the monopolist in this case, is given by the following specialization of equation (15):[14]

$$\begin{aligned} E\Pi(u) &= \Theta_{12}(\pi_1^{1st} - u) + \alpha(u)\Theta_{12}\pi_2^{1st} \\ &\quad + [2 - \Theta_{12} - \alpha(u)\Theta_{12}]\Big\{v_2 q_2(2, \alpha(u)) - c\Big(q_2(2, \alpha(u))\Big)\Big\}\,, \qquad (21) \end{aligned}$$

where we have used the fact that $u = (v_1 - v_2)q_2(2, \alpha(u))$ and also substituted q_2^{1st} for $q_2(1, \alpha)$. Since the functions $\alpha(u)$ and $q_2(2, \alpha(u))$ are continuous functions of u on the closed interval $[0, u_{max}]$, $E\Pi(u)$ is a bounded, continuous function of u which has a maximum on this interval.

Let u^* denote a value of u that maximizes $E\Pi(u)$ on the interval $[0, u_{max}]$. The equilibrium strategy for the monopolist in the first period is to offer the single item: $(p_1^{1st} - u^*,\ q_1^{1st})$.

We now consider the monopolist's second-period strategy. The monopolist observes what he offers in the first period and what, if any, items were purchased. Based on these observations and his knowledge of the buyers' equilibrium strategies, he revises his beliefs about the probability that there is a type 1 buyer remaining. In particular, the seller observes in the first period either zero, one, or two purchases from either a type I or type II menu. If two purchases occur, then the seller has no choices left to make. If one purchase occurs, he can distinguish whether or not the purchased item is the best item for a type 1 buyer. The beliefs and second-period actions of the seller are summarized in Table 1.

[Table 1 goes here]

The seller's beliefs follow from Bayes' rule when it applies. As Table 1 indicates, there are five cases where Bayes' rule applies. When the best item for a type 1 buyer is selected from a type I menu and $0 \leq u \leq u_{max}$, the seller realizes that this would occur in equilibrium only when a type 1 buyer makes the purchase; he therefore believes that no type 1 buyer remains in the final period. When no purchase occurs in the same circumstance, the seller realizes that—even with a type 1 buyer present—this happens with probability $1 - \alpha(u)$; he therefore concludes that there is probability $\Theta_{12}(\alpha(u))$ that one of the remaining buyers is type 1. If the seller offered a type

Table 1. Seller's Second-Period Beliefs and Strategy

History	Probability that a type 1 remains	Action
Type II menu and		
1 purchase	$1/2\ \Theta_{12}$	p_1^{1st}, q_1^{1st}
0 purchase	Θ_{12}	p_1^{1st}, q_1^{1st}
Type I menu with $u \in [0, u_{max}]$ and		
1 purchase (not best for type 1)	$1/2\ \Theta_{12}$	p_1^{1st}, q_1^{1st}
1 purchase (best for type 1)	0 ◇	p_2^{1st}, q_2^{1st}
0 purchase	$\Theta_{12}(\alpha(u))$ ◇	optimal menu†
Type I menu with $u > u_{max}$ and		
1 purchase (not best for type 1)	$1/2\ \Theta_{12}$	p_1^{1st}, q_1^{1st}
1 purchase (best for type 1)	0 ◇	p_2^{1st}, q_2^{1st}
0 purchase	0 ◇	p_2^{1st}, q_2^{1st}
Type I menu with $u < 0$ and		
1 purchase	$1/2\ \Theta_{12}$	p_1^{1st}, q_1^{1st}
0 purchase	Θ_{12} ◇	p_1^{1st}, q_1^{1st}

◇ Beliefs dictated by Bayes' rule

† See text

I menu but mistakenly made it worth $u > u_{max}$ to a type 1 buyer and if the best item for such a buyer was purchased, then the seller concludes that no type 1 buyer remains in the second period. The seller also infers that no type 1 buyer is present after observing zero purchases from such a menu. Finally, if the seller offered a type I menu but mistakenly made it worth $u < 0$ to a type 1 buyer, then no purchases should occur; if none does, the seller would retain his prior beliefs about the presence of a type 1 buyer.

In the other five cases listed in Table 1, the seller observes behavior which is inconsistent with the buyers' equilibrium strategies. For simplicity, we assume in all such cases that the seller's expectations about the types of the remaining buyers are formed using his prior beliefs. In particular, he believes when two buyers remain that the probability that one of them is a type 1 is still Θ_{12}. If only one buyer remains, the seller assigns probability $(1/2)\Theta_{12}$ to the event that the buyer is a type 1.

We now verify that the proposed strategy for each player is optimal at each stage. Since buyers of either type who purchase nothing receive zero utility, it is optimal for anyone remaining in the last period—no matter what his beliefs—to purchase the best item if it yields nonnegative utility. A type 1 buyer can anticipate receiving only zero utility in the second period whenever the first-period menu is type II. Faced with such a menu, it is therefore also optimal for him to accept the best item in the first period that yields nonnegative utility. For the same reason, it is always optimal for a type 2 buyer to follow this strategy in the first period.

Now consider a first-period type I menu that provides maximum utility u to a type 1 buyer. If $u < 0$, it is clearly optimal for the buyer to reject such an offer. If $u > u_{max}$, then it is optimal for a type 1 buyer to accept the best item from this menu since if he rejects it he will get utility equal to u_{max} in the second period.

A type 1 buyer who in the first period rejects a type I offer that provides $0 \leq u \leq u_{max}$ receives utility $(v_1 - v_2)q_2(2, \alpha(u))$ in the second period. From the definition of $\alpha(u)$ it follows that $u = (v_1 - v_2)q_2(2, \alpha(u))$, so that a type 1 buyer is indifferent between rejecting and accepting such a first-period offer. In particular, it is optimal for a type 1 buyer to accept the best item from such a menu with probability $\alpha(u)$.

Having shown that the strategy of each type of buyer is sequentially rational, we consider the strategy of the seller. In two of the cases listed in Table 1, two buyers remain and the seller retains his prior probability (Θ_{12}) that one of them is type 1. Since the prior probability distribution gives rise to a corner in a one-period game, it is optimal for the seller to offer (p_1^{1st}, q_1^{1st}) in the second period. In four other cases listed in Table 1, one buyer remains and the seller assesses the probability that he is type 1 at $(1/2)\Theta_{12}$. In this circumstance, it is optimal for the monopolist to offer the single item (p_1^{1st}, q_1^{1st}) in the second period for the same reason that this offer is optimal in the static Mussa-Rosen game.[15] In three of the cases listed in Table 1, the seller assigns zero probability to the event that any remaining buyer is type 1. In this circumstance, it is optimal to offer the single item (p_2^{1st}, q_2^{1st}).

In the remaining case in Table 1, the first-period menu is type I and although

$u \in [0, u_{max}]$, no purchases occur. Since the monopolist's updated beliefs are specified by $\Theta_{12}(\alpha)$, it is optimal for the monopolist in the second period to make the Mussa-Rosen offer determined by these beliefs.

It remains to show that the seller's proposed strategy in the first period is optimal. The monopolist can achieve W^{MR} by offering the best type II menu initially or, alternatively, by offering a type I menu with $u < 0$. However, these alternatives are strictly dominated by offering a type I menu with $u = 0$. Suppose the monopolist offered the single item $(p_1^{1st},\ q_1^{1st})$. Even this (possibly suboptimal) offer yields the seller an expected profit $E\Pi(0)$ which strictly exceeds W^{MR}. Using equation (21) and the fact that $q_2(2, \alpha_0) = 0$, we obtain the following expression:

$$E\Pi(0) = \Theta_{12}\pi_1^{1st} + \alpha_0\Theta_{12}\pi_2^{1st} = W^{MR} + \alpha_0\Theta_{12}\pi_2^{1st}\,. \tag{22}$$

Type I menus with $u > u_{max}$ are strictly dominated by the single-item menu $(p_1^{1st} - u_{max},\ q_1^{1st})$ since the latter would also always be accepted and gives away strictly less surplus.

Of all items that provide utility u to a type 1 buyer, the item $(p_1^{1st} - u,\ q_1^{1st})$ provides the greatest profit for the seller. Together with the previous remarks, this implies that the optimal initial offer of the monopolist is a type I menu containing only a single item with quality q_1^{1st} offered at a price $p_1^{1st} - u$ for some $u \in [0, u_{max}]$. In particular, the proposed strategy of offering the single item $(p_1^{1st} - u^*,\ q_1^{1st})$ must be optimal since u^* was chosen to maximize the seller's expected profit, $E\Pi(u)$.

This concludes the demonstration that the specified strategies and beliefs constitute a perfect Bayesian equilibrium.

Since $E\Pi(u^*) \geq E\Pi(0) > W^{MR}$, the monopolist in this example can always increase his profits relative to the static Mussa-Rosen case by making a sequence of offers. However, he will not generally be able to capture the entire surplus. It is either optimal for the monopolist to leave a type 1 buyer with utility $u^* > 0$ or to make the offer $(p_1^{1st},\ q_1^{1st})$ in the first period. In the latter case, equation (22) indicates that the monopolist's expected profit will be less than the total expected surplus $(\ \Theta_{12}\pi_1^{1st} + [2 - \Theta_{12}]\pi_2^{1st}\ )$.

The above example is a natural extension to incomplete information of the analysis of the "corner" parameter regime in section 2. Equation (20) implies that as Θ_{12} approaches 1, α_0 also approaches 1. Together with equation (22), this implies that $E\Pi(0)$ approaches $\pi_1^{1st} + \pi_2^{1st}$ as Θ_{12} approaches 1. Hence, the initial offer by the seller which was exactly optimal under complete information is approximately optimal (since it extracts almost the entire expected surplus) in this example as the uncertainty about the aggregate population of buyers becomes small. Moreover, as in the "corner" parameter regime with complete information and a discount factor which approached one, the monopolist with incomplete information (and $\beta = 1$) obtains approximately the surplus of a perfectly discriminating monopolist as Θ_{12} approaches 1.

Taken together, the results in this section indicate that correlation — either positive or negative — is needed in the equilibria we consider if the seller is to benefit from the opportunity to make a sequence of offers. Sequencing enables the seller to make an offer to one buyer conditional on the *prior* response of the other buyer even though, in contrast to a static direct revelation game, the seller is assumed to be unable to condition his offers to one buyer on the *contemporaneous* responses of the other buyer. Our second example with correlated types shows that the inability to condition on contemporaneous behavior can prevent the monopolist from using a sequence of offers to extract the full surplus from buyers.

In our two examples, the seller is uncertain about the realized valuations of the buyers but no type 1 buyer is uncertain of the valuation of the other buyer. This characteristic of our examples shortens the exposition but is not responsible for our results. Given more space, examples could be described where *all* players are uncertain and yet a sequence of offers leads to higher expected profits than a single offer; in some of these examples the seller approximates the expected profits of perfect price discrimination.

4. Concluding Remarks

The model in this paper is related to the one in Bagnoli et al. (1989) although our focus there was entirely different. In the previous paper, we assumed that the monopolist could not alter the quality of the durables he produced. Hence, every item offered at a point in time had the same price, and inducing point-in-time self-selection was infeasible. In addition, we assumed the seller had complete information about the valuation of each buyer. Our focus in that paper was instead on Coase's famous conjecture that in the continuous-time limit, buyers always receive the entire social surplus.

The current paper contributes to the literature on second-degree price discrimination. That literature has focused on the benefits to a monopolist of inducing buyers to self-select at a point in time by offering a single menu of items with different characteristics. As our paper shows, when the monopolist is allowed to offer a sequence of menus over time, the opportunity to induce intertemporal self-selection provides the seller—whether he has complete or incomplete information—with a potent additional means of extracting surplus when the discount factor is sufficiently high. In many of the examples we consider, it is optimal for the monopolist to abandon multi-item menus altogether and rely solely on intertemporal self-selection (except perhaps in the last period, if one exists).

Our paper also provides a connection between the literatures on sequential bargaining and mechanism design. The ability to make a sequence of offers provides the seller with a degree of flexibility in pricing greater than that which is usually assumed in static models of price discrimination but less than is usually assumed in the literature on mechanism design.

References

Arrow, K. (1986) "Agency and the Market," chapter 23 in *The Handbook of Mathematical Economics, Vol. 3*, edited by K. Arrow and M. Intrilligator, North-Holland, pp. 1183–1195.

Bagnoli, M., S.W. Salant, and J. Swierzbinski (1989) "Durable-Goods Monopoly with Discrete Demand," *Journal of Political Economy* **97**, pp. 1459–1478.

Bagnoli, M., S.W. Salant, and J. Swierzbinski (1990) "Price Discrimination and Intertemporal Self-Selection," Center for Research on Economic and Social Theory (CREST) Working Paper, University of Michigan.

Bulow, J. (1982) "Durable-Goods Monopolists," *Journal of Political Economy* **90**, pp. 314–332.

Coase, R. (1972) "Durability and Monopoly," *Journal of Law and Economics* **15**, pp. 143–149.

Cremer, J. and R. McLean (1985) "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica* **53**, pp. 345–361.

Demski, J. and D. Sappington (1984) "Optimal Incentive Contracts with Multiple Agents," *Journal of Economic Theory* **33**, pp. 152–171.

Fudenberg, D., D. Levine, and J. Tirole (1985) "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game-Theoretic Models of Bargaining*, edited by A. Roth, Cambridge University Press, pp. 73–98.

Gul, F., H. Sonnenschein, and R. Wilson (1986) "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory* **39**, pp. 155–190.

Holmstrom, B. (1982) "Moral Hazard in Teams," *Bell Journal of Economics* **13**, pp. 324–340.

Kahn, C. (1986) "The Durable Goods Monopolist and Consistency with Increasing Costs," *Econometrica* **54**, pp. 275–294.

Mussa, M. and S. Rosen (1978) "Monopoly and Product Quality," *Journal of Economic Theory* **18**, pp. 301–317.

Sobel, J. and I. Takahashi (1983) "A Multistage Model of Bargaining," *Review of Economic Studies* **50**, pp. 411–426.

Stokey, N. (1981) "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics* **12** pp. 112–128.

Tirole, J. (1988) *The Theory of Industrial Organization*, MIT Press, Cambridge, Mass.

Wang, G. H. (1991) "Multi-Dimensional Sequential Bargaining Under Incomplete Information," draft working paper.

Footnotes

1. Contributions to this literature include Bagnoli, Salant, and Swierzbinski (1989), Bulow (1982), Coase (1972), Fudenberg, Levine, and Tirole (1985), Gul, Sonnenschein and Wilson (1986), Kahn (1986), Sobel and Takahashi (1983) and Stokey (1981).

2. The appendix of Bagnoli et al. (1990) extends our two-buyer, infinite-horizon analysis to the case of a complete-information, infinite-horizon game where there are initially a finite number of buyers of each of a finite number of types.

3. See Fudenberg, Levine, and Tirole (1985) for an example of a bargaining game where a seller who can offer only a fixed quality is disadvantaged by the ability to make an infinite sequence of offers rather than a single take-it-or-leave-it offer. Wang (1991) studies both finite- and infinite-horizon single-buyer games where a monopolist can offer menus of items with different qualities and concludes that multiple offers provide no advantage.

4. The assumption that $v_2 > c'(0)$ is sufficient to ensure that both buyers are served under perfect price discrimination and that at least buyer 1 is served in Mussa-Rosen's problem.

5. A trivial multiplicity of equilibria arises because the seller can always add items to the menu which are unacceptable to every remaining buyer. Even in the absence of such items, a second trivial multiplicity of equilibria arises because of buyer 2's indifference about accepting offers off the equilibrium path yielding zero utility. However, the play of the game and the payoffs which result in these equilibria are unique.

6. In the first period of the game, the expected numbers of buyers of each type are related to the prior probabilities via the equations:

$$\begin{aligned} \bar{n}_1 &= 2\Theta_{11} + \Theta_{12} \\ \bar{n}_2 &= 2\Theta_{22} + \Theta_{12} \, . \end{aligned}$$

7. See Tirole (1988) for a discussion of the concept of perfect Bayesian equilibrium.

8. For brevity, we will also refer to the single-item menu (p_2^{1st}, q_2^{1st}) which is optimal for the monopolist when he believes that all remaining buyers are type 2 as a Mussa-Rosen menu.

9. Let $r(2)$ denote the ratio of these expected values when two buyers remain in the second period and $r(1)$ denote their ratio when only one buyer remains. Bayes' rule implies that $r(2) = \frac{2(1-\alpha)^2\Theta_{11}+(1-\alpha)\Theta_{12}}{2\Theta_{22}+(1-\alpha)\Theta_{12}}$ and $r(1) = \frac{2(1-\alpha)\Theta_{11}}{\Theta_{12}}$. Since independence implies that $\Theta_{12}^2 = 4\Theta_{11}\Theta_{22}$, the reader can verify that $r(1) = r(2)$.

10. To define correlation precisely, arbitrarily label the two buyers A and B and define two associated random variables as follows. If A is a type 1, the associated random variable has the value 1; if A is a type 2, the random variable has the value 2. Define the second random variable analogously to reflect the type of buyer B. The

correlation coefficient between these two random variables (ρ) can now be defined in the standard way. We obtain the result:

$$\rho = \frac{\Theta_{11} - (\Theta_{11} + .5\Theta_{12})^2}{(\Theta_{11} + .5\Theta_{12}) - (\Theta_{11} + .5\Theta_{12})^2} .$$

In our first example, the correlation coefficient is positive ($\rho = 1$). In the second example, any negative correlation coefficient can be produced ($-1 \leq \rho < 0$) by a suitable choice of Θ_{12}.

11. For the two-buyer, single-period game with $\Theta_{11} = 0$, it is the case that $\bar{n}_1 = \Theta_{12}$ and $\bar{n}_2 = 2 - \Theta_{12}$. Moreover, when equations (11) and (12) are substituted into equation (13), the Kuhn-Tucker conditions for the expected-profit-maximizing qualities imply that $q_2 = 0$ is optimal if and only if:

$$\bar{n}_2 v_2 - \bar{n}_1(v_1 - v_2) - \bar{n}_2 c'(0) \leq 0 .$$

Substituting for $\bar{n}_1$ and $\bar{n}_2$ in the above inequality and collecting terms results in the inequality in the text.

12. For the parameter values in our second example, there is no perfect Bayesian equilibrium where the type 1 buyer uses a pure strategy after every first-period offer. For suppose that the monopolist offers a type I menu in the first period that provides $0 < u < u_{max}$ to a type 1 buyer. If the type 1 buyer's strategy is to accept such an offer, then after zero acceptances are observed, Bayes' rule implies that both buyers must be type 2. In this case, the monopolist optimally offers the single item (p_2^{1st}, q_2^{1st}) in the second period and the type 1 buyer could have obtained greater utility by deviating from the proposed equilibrium strategy and rejecting the initial offer. On the other hand, suppose that the type 1 buyer's strategy is to reject the previously mentioned first-period offer. In this case, following zero acceptances, Bayes' rule specifies that the monopolist's posterior beliefs must be the same as his prior beliefs. With these beliefs, the monopolist optimally offers the single item (p_1^{1st}, q_1^{1st}) in the second period and the type 1 buyer could have obtained greater utility by accepting the initial offer.

13. The definition of $q_2(2, \alpha)$ implies that $q_2(2, \alpha_0) = 0$ and $q_2(2, 1) = q_2^{1st}$. Moreover, in the two-buyer game for which $q_2(2, \alpha)$ is defined, the expected number of type 1 buyers is $\bar{n}_1 = \Theta_{12}(\alpha)$ and the expected number of type 2 buyers is $\bar{n}_2 = 2 - \Theta_{12}(\alpha)$. Given our assumptions about $c(q)$, the Kuhn-Tucker conditions imply that $q_2(2, \alpha)$ is a continuous, strictly increasing function of α on the interval $[\alpha_0, 1]$. Hence, $q_2(2, \alpha)$ has a continuous, strictly increasing inverse on the interval $[0, q_2^{1st}]$. Call this inverse $h(x)$. Then $\alpha(u) = h(u/(v_1 - v_2))$.

14. For $0 \leq u \leq u_{max}$, it is straightforward to verify that the item ($p_1^{1st} - u$, q_1^{1st}) provides negative utility for a type 2 buyer and so would be rejected by such a buyer. The maximum price that a type 2 buyer would be willing to pay for an item with quality q_1^{1st} is $v_2 q_1^{1st}$. Using the fact that $u_{max} = (v_1 - v_2)q_2^{1st}$, it is easy to show that $p_1^{1st} - u_{max} > v_2 q_1^{1st}$.

15. In a game with a single buyer where the probability that the buyer is type 1 is $(1/2)\Theta_{12}$, the expected number of type 1 buyers is $\bar{n}_1 = (1/2)\Theta_{12}$ and the expected number of type two buyers is $\bar{n}_2 = 1 - (1/2)\Theta_{12}$. Hence, the ratio $\bar{n}_1/\bar{n}_2$ is the same as in the original two-buyer game. As we have already observed, the Mussa-Rosen offer (with one or more buyers) depends on the monopolist's beliefs about the buyers' types only through this ratio.

Recent Crest Working Papers

91–01: Mark Bagnoli, Stephen W. Salant, Joseph E. Swierzbinski, "Price Discrimination and Intertemporal Self-Selection," June, 1990.

91–02: Howard Doran, Jan Kmenta, "Multiple Minima in the Estimation of Models with Autoregressive Disturbances," April 1991.

91–03: Ted Bergstrom, Mark Bagnoli, "Courtship as a Waiting Game," April 3, 1991.

91–04: Ted Bergstrom, Jeffrey K. MacKie-Mason, "Peak-Load Pricing-With and Without Constrained Rate of Return," April 2, 1991.

91–05: Laura E. Kodres, Daniel P. O'Brien, "The Existence of Pareto Superior Price Limits and Trading Halts," February 1991.

91–06: Theodore Bergstrom, David Lam, "The Effects of Cohort Size on Marriage Markets in Twentieth Century Sweden," November 1989.

91–07: Theodore Bergstrom, David Lam, "The Two-Sex Problem and the Marriage Squeeze in an Equilibrium Model of Marriage Markets," April 1991.

91–08: Greg Shaffer, "Capturing Strategic Rent: Full-line Forcing, Maximum Resale Price Maintenance, Brand Discounts and Aggregate Rebates," March 1991.

91–09: Jeffrey K. MacKie-Mason, Roger H. Gordon, " Taxes and the Choice of Organizational Form," April 3, 1991.

91–10: Hal R. Varian, "A Solution to the Problem of Externalities When Agents are Well-Informed," September 1991.

91–11: Hal R. Varian, "The Economics of User Interfaces," September 1991.

91–12: Eduardo Ley, Hal R. Varian, "A Note on the Dow–Jones Index," September 1991.

91–13: Hal R. Varian, "Goodness of Fit for Revealed Preference Tests," September 1991.

91–14: Hal R. Varian, "Sequential Provision of Public Good," September 1991.

91–15: Ken Binmore, "A Liberal Leviathan," September 1991.

91–16: Ken Binmore, "DeBayesing Game Theory," September 1991.

91–17: Ken Binmore, "Bargaining and Morality," September 1991.

91–18: Ted Bergstrom, "When Non–Transitive Relations Take Maxima and Competitive Equilibria Can't be Beat."

92–01: Lung-fei Lee, "Asymptotic Distribution of the Maximum Likelihood Estimator for a Stochastic Frontier Function Model with a Singular Information Matrix," May, 1992.

92–02: Stephen Salant, Karen Kalat, Ana Wheatcroft, "Surviving Winter: A Fitness-Based Explanation of Hoarding and Hibernation," February, 1992.

92–03: Gérard Gaudet, Stephen Salant, "The Limits of Monopolization Through Acquisition: Further Results," February, 1992.

92–04: Jonathan Cave, Stephen W. Salant, "Cartel Quotas Under Majority Rule," February, 1992.

92–05: Mark Bagnoli, Stephen W. Salant, Joseph E. Swierzbinski, "Intertemporal Self-Selection with Multiple Buyers Under Complete and Incomplete Information," April, 1992.

92–06: Daniel P. O'Brien, Greg Shaffer, "Non-linear Contracts, Foreclosure, and Exclusive Dealing," April, 1992.

92–07: Lung-fei Lee, Jungsywan H. Sepanski, "Consistent Estimation of Linear and Nonlinear Errors-in-Variables Models with Validation Information," April, 1992.

92–08: Lung-fei Lee, "Semiparametric Minimum-distance Estimation," November, 1991.

92–09: Michelle J. White, "Corporate Bankruptcy as a Filtering Device," March 29, 1992.

92–10: Ted Bergstrom, Robert Schoeni, "Income Prospects and Age at Marriage," March 4, 1992.

92–11: James Andreoni, Ted Bergstrom, "Do Government Subsidies Increase the Private Supply of Public Goods?," February 6, 1992.