

**Nonlinear Price Discrimination
with a Finite Number of Consumers
and Constrained Recontracting**

Jonathan Hamilton

and

Steven Slutsky

Department of Economics
Warrington College of Business Administration
University of Florida
Gainesville FL 32611-7140 USA

March 2004

An earlier version of this paper was presented at the 1999 Southeastern Economic Theory Meeting, the 2000 Conference on Public Economics at Academia Sinica, and seminars at Florida International University, Sungkyunkwan University, the University of Copenhagen, the University of British Columbia, and the University of Virginia. We thank an anonymous referee, Michael Blake, Jacques Crémer, Tracy Lewis, Richard Romano, David Sappington, Art Snow, and Barbara Spencer for helpful comments. We thank the Warrington College of Business Administration for financial support for this research. All errors are solely our responsibility.

Abstract

In the monopoly nonlinear pricing problem with unobservable consumer types, we study mechanisms in which the firm makes the set of consumer options conditional on the aggregate reports of consumer types it receives. Previous mechanisms that exploit knowledge of the true type distribution often have multiple equilibria or use noncredible contracts off the equilibrium path. When the monopolist can replace contracts after initial reports subject to the constraint that truthful consumers are not made worse off, the outcome is essentially the same as when the monopolist has full information. This holds whether or not the monopolist can make offers to consumers who reject all original contract offers. When the monopolist must guarantee nonnegative surplus to all truthful consumers in all contingencies, the equilibrium outcome has undistorted contracts but lower profits for the monopolist.

1. Introduction

The best modern interpretation of second-degree price discrimination (or declining-block pricing) is that it arises as the solution to an asymmetric information problem. The monopolist knows the distribution of consumer types but cannot identify demand curves of individual consumers. While the monopolist takes the distribution of consumer types into consideration in choosing the optimal schedules, it does not use this information about market aggregates as aggressively as possible. In particular, the conventional solution has the monopolist choosing the nonlinear pricing schedule subject to incentive compatibility constraints (so that each consumer would truthfully reveal his type to the monopolist), but the contracts offered do not depend on the number of consumers who request each of the contract choices. One justification for this approach is that the monopolist is assumed to be selling to a large (infinite) number of consumers.¹

An alternative assumption that we explore in this paper is that the monopolist has a small (finite) set of potential customers and knows the exact distribution of types. With precise information on the number of each type of consumer in the market, the monopolist can detect that some consumers seek to pay a lower average price per unit by reporting that they prefer a smaller quantity than they actually do. Using this information significantly changes the nature of the monopoly's optimal pricing policies.

Each of these specifications is most appropriate in different contexts. The standard model may best fit a final good monopolist selling to a large number of individuals. Even though this number is still finite, uncertainty about the exact distribution of types can smooth things out, making a model with a continuum of consumers appropriate. On the other hand,

¹ With a continuous distribution of types, Hammond [1979] shows that one cannot improve on such a solution by conditioning one types payments on the reports of other types.

an intermediate good monopolist may sell to a small number of oligopolists and, through access to general market research, may have reasonably precise information about the distribution of types to whom it is selling.

Previous models have exploited the seller's information about consumer aggregates in a variety of ways. In a perfect information setting, Levine and Pesendorfer [1995] allow a monopolist to precommit to a pricing strategy in which no sales take place if any consumer chooses to purchase less than the quantity which yields zero surplus to the consumer. They also show that such an equilibrium is not robust to the introduction of noise in players' actions.

Bagnoli, Salant, and Swierzbinski [1989] consider a durable-good monopoly and find that the monopolist can extract all surplus by making a sequence of price offers which depend on the history of purchases. This equilibrium is subgame perfect, but not unique.² Bagnoli, Salant, and Swierzbinski [1995] study a similar mechanism in a model of quality differences and find the monopolist can extract all surplus. We discuss the differences between their mechanism and ours in the conclusion.

With asymmetric information, the idea that correlation between consumer types may help a monopolist extract additional surplus from consumers has been studied elsewhere, especially in auctions. In fact, Crémer and McLean [1985, 1988], Brusco [1998], and Spiegel and Wilkie [2000] develop mechanisms by which the seller can extract the full surplus under certain conditions. Fudenberg and Tirole [1991, p. 294-295] discuss the fact that these mechanisms may require large payments by consumers in certain outcomes of the game. Our correlation among types is a special case of the distributions studied by Crémer and McLean

² While other equilibria are not considered in the paper, it is clear that Coase's solution remains an equilibrium of the game.

because the aggregate distribution is common knowledge among consumers and the firm. However, our mechanism bounds the losses to agents in outcomes off the equilibrium path and some of our assumptions are more general. We discuss differences between these models and ours more fully in the final section.

The asymmetric information monopoly problem bears a strong analogy to the optimal income tax problem with a finite number of consumers.³ Piketty [1993] shows that the government can implement any allocation on the full-information Pareto frontier as a dominance-solvable equilibrium by allowing the government to condition individual tax schedules on the aggregate reports from all workers. Thus, in contrast to the models of Mirrlees [1971] with a continuum of ability levels and Stiglitz [1982] with two ability levels and a continuum of workers of each type, no worker faces a positive marginal tax rate. Piketty calls such mechanisms “generalized tax schedules.” In effect, whenever too many taxpayers claim to be of low ability, the bundle given to these taxpayers is worse than the bundle intended for low ability taxpayers when everyone reports abilities truthfully.

While researchers prior to Piketty were aware that mechanisms using aggregate information could sustain efficient allocations that violated the conventional incentive compatibility conditions, they generally disregarded such solutions since they viewed them as being plagued by problems of multiple equilibria. They thought the outcome when even one high-ability taxpayer claimed to be of low ability had to be set as very bad for low-ability types. In that case, a low-ability type would be better off to claim to be of high ability to

³ Similarities between the monopoly problem and the optimal tax problem have long been recognized. For example, Stiglitz [1977] observes that one is the dual of the other, and Maskin and Riley [1984] show strong similarities in the characterizations of the optimal solutions in the two cases.

avoid the punishment outcome.⁴ Although Piketty circumvents this difficulty and obtains a unique equilibrium, his approach has other unsatisfactory properties. First, the bundles offered taxpayers off the equilibrium path (after at least one taxpayer has claimed to be different from his true type) do not necessarily satisfy the government's budget constraint and hence may not be feasible. Second, *ex post*, the bundles given to consumers are not necessarily the optimal choices for the government given the information it has about taxpayers at that point. Hamilton and Slutsky [2003] find that imposing budget balance on and off the equilibrium path and requiring the allocations off the equilibrium path to satisfy certain optimality conditions limits the set of efficient allocations that the mechanism designer can achieve.

The analogies between the nonlinear pricing and optimal income tax problems are significantly weakened by two differences. First, the firm does not face a budget constraint, in contrast to a redistributive government. Thus, budget balance off the equilibrium path does not affect the monopolist's problem. However, the problem remains of whether the bundles chosen off the equilibrium path are optimal given the monopolist's information at that point. A monopolist, having announced quantity-outlay pairs conditional on the number of consumers who request each of the pairs offered, may want to revise these offers after consumers make their choices if misrevelation has occurred. The off-equilibrium-path pairs were initially chosen to sustain the truthful outcome and not to maximize profits conditional on a false report.

Second, a monopolist faces participation constraints that customers may choose to buy nothing if offered a very bad deal, in contrast to the fixed populations in the optimal income

⁴ Such problems would arise in a variant of the Levine-Pesendorfer mechanism with multiple consumer types.

tax model.⁵ One modeling issue that arises in this context is whether the participation constraints apply only on the equilibrium path or off the equilibrium path as well. We consider both cases in our analysis.

If the monopolist does not commit to fulfill contracts initially offered, consumers may be unwilling to reveal their types to the monopolist. The possibility we analyze in this paper is that the monopolist may be able to commit to fulfill the contracts offered, but it may not be able to commit not to revise those contracts with consent from consumers. Legal action by consumers may prevent the monopolist from making consumers who reported their types truthfully worse off than promised. However, even legally binding contracts can be altered if neither party is harmed. The monopolist is able to replace contracts, so long as the new contracts do not lower the utility of someone who has revealed her type truthfully.⁶ That is, a policy change which affects those who reveal themselves to be of a certain type could not lower the utility of someone of that type. However, the policy change could harm someone who has misrevealed, since an individual without “clean hands” would have his case thrown out of court. In this respect, our partial commitment differs in spirit from renegotiation in games (see, for example, Farrell and Maskin [1989]) and in principal-agent problems (see, for example, Fudenberg and Tirole [1990]), where renegotiation must lead to a Pareto improvement.⁷ Baron and Besanko [1987] analyze a principal-agent model in which the principal can guarantee only that any truth-telling agent will earn nonnegative profit in the

⁵ Issues of potential migration are usually ignored in optimal income tax problems.

⁶ An alternative to replacing contracts would be to let the monopolist offer supplemental quantity-outlay pairs to consumers who had purchased a particular bundle initially. Such a mechanism permits the same type of contract revisions as we model here.

⁷ Thus, after some misrevelation, renegotiation would require making consumers who told the truth, consumers who misreported their types, and the monopolist no worse off.

second period of a regulatory regime. Our partial commitment is similar in spirit, except that we allow the principal to guarantee any feasible utility level to truth-telling agents (and to condition the levels of these guarantees on the aggregate reports).⁸

We also require that the monopolist must induce a game among consumers which has a unique equilibrium. In the literature on implementation, several approaches exist to ensure uniqueness. When the outcome is only a Nash equilibrium, uniqueness does not often result, while requiring truth to be a dominant strategy for all consumers would essentially guarantee uniqueness. Mookerjee and Reichelstein [1992] consider when the requirement of having dominant strategies can be imposed with no loss. The monotonicity condition they find is not satisfied in our context. Thus we take an alternative approach to ensuring uniqueness by requiring a weaker condition than existence of dominant strategies, namely that the game be dominance solvable (solvable by iterated deletion of strictly dominated strategies). This is the approach used by Piketty [1993]. It is also used in the literature on virtual implementation as in Abreu and Sen [1991], Abreu and Matsushima [1992a, 1992b], and Duggan [1997]. These papers show that a planner can implement an allocation arbitrarily close to any outcome in many environments by use of a lottery in which the desired outcome is a prize with high probability. When the monopolist can commit to contracts, our model is consistent with the virtual implementation framework. However, we do not use lotteries, but implement an allocation close to the full-surplus extraction outcome with a relatively simple mechanism.⁹

⁸ The idea that sellers' offers to consumers are only a one-sided commitment because a contract would not prevent the firm from revising its offer in a way that makes the consumer better off has been studied in duopoly models [Hamilton, MacLeod, and Thisse, 1991].

⁹ We do not need lotteries because we have a continuum of alternatives since the numeraire is divisible.

In the simplest case, even though there are many information sets, the monopolist need only offer two different pairs of bundles.

We move outside the implementation framework with our main results when the monopolist cannot commit and may revise contracts subject to partial commitment constraints. Now, the monopolist cannot be viewed as a “mechanism designer” operating outside the game, but must be viewed as a player in the game. As part of the overall game, the monopolist specifies a game among consumers by offering a menu of contracts and revising them after consumers reveal their types.

We solve for the sequential equilibria of the nonlinear pricing games with this form of updating contracts. With contract revision subject to partial commitment, both consumer types purchase the efficient quantities, and the monopolist still extracts almost the entire surplus. This also holds if the monopolist can make offers to consumers who initially opted out. Restrictions on the minimum utility offered to truthful consumers both on and off the equilibrium path limit the monopolist’s profit but outcomes are still efficient.

Section 2 presents the basic model with two consumer types. Section 3 analyzes the contract revision game and Section 4 analyzes the game with offers to those who opt out in the first stage. Section 5 considers contract limitations. Section 6 considers some additional applications of our approach when the monopolist has less precise information. In Section 7, we present our conclusions and compare our results to other literature.

2. The Nonlinear Pricing Model with Asymmetric Information

We study a simple model of discriminatory pricing with asymmetric information. There are two types of consumers whose preferences over the quantity of the monopolist’s output and the amount they pay can be described by the continuous, strictly quasiconcave payoff

functions: $V^1(q, R)$ and $V^2(q, R)$, where q is the quantity of the monopolist's good and R is the payment to the firm. Thus, $\partial V^i / \partial q > 0$ and $\partial V^i / \partial R < 0$. We make the standard single-crossing assumption that a type 2 indifference curve in (q, R) space is steeper than a type 1 indifference curve at any point of intersection (type 2 has a greater willingness to pay for the good than type 1 at any price). See Figure 1. We normalize each type's utility function so that $V^i(0,0) = 0$. There are N consumers of type 1 and $T - N$ consumers of type 2. For simplicity we assume that the monopolist has a constant marginal cost of production.

Let us first consider the full-information problem. Let contracts $\hat{\alpha}$ and $\hat{\beta}$ denote the (q, R) pairs that are the solution to:

$$\begin{aligned} \Pi^F \equiv \quad & \text{Max}_{q^1, q^2, R^1, R^2} \quad NR^1 + (T - N)R^2 - c(Nq^1 + (T - N)q^2) \\ & \text{s.t.} \quad V^1(q^1, R^1) \geq V^1(0, 0) = 0 \\ & \quad \quad V^2(q^2, R^2) \geq V^2(0, 0) = 0. \end{aligned}$$

The constraints are the individual rationality (IR) constraints that require that each type does at least as well by buying from the monopolist as by purchasing nothing. Because of constant marginal costs, in this problem, the monopolist is actually dealing with two completely independent markets. Contract $\hat{\alpha}$ maximizes profit from type 1 consumers subject to the individual rationality constraint for these consumers. It will have marginal willingness to pay equal to marginal cost and the payment will extract all surplus so type 1's (IR) constraint holds with equality. We assume that all consumers are willing to pay more than marginal cost for at least some quantity of the good, so that $\hat{\alpha} \neq (0, 0)$.

Similarly, contract $\hat{\beta}$ maximizes profit from type 2 consumers subject to the individual rationality constraint for these consumers. The same properties with respect to type 2's must hold. The $\hat{\beta}$ contract would yield negative surplus for type 1 consumers so $V^1(\hat{\alpha}) = V^1(0, 0)$

$> V^1(\hat{\beta})$. Clearly, a type 2 consumer envies the contract offered to a type 1 so $V^2(\hat{\alpha}) > V^2(\hat{\beta}) = V^2(0, 0)$. See Figure 1.

The conventional problem facing the monopolist under asymmetric information is to maximize profits by offering a pair of quantity-outlay (q, R) contracts that satisfy incentive compatibility (IC) constraints. The monopolist's offer to consumers does not depend on how many consumers choose each contract. If more consumers choose one contract than the monopolist anticipated, it is unable to revise the contract. The monopolist anticipates that consumers will reveal their types truthfully and the offers are the same both on and off the equilibrium path. When the monopolist chooses to sell to everyone, we can write the profit maximization problem with asymmetric information as:

$$\begin{aligned} \Pi^B \equiv \quad & \text{Max}_{q^1, q^2, R^1, R^2} \quad NR^1 + (T - N)R^2 - c(Nq^1 + (T - N)q^2) \\ & \text{s.t.} \quad V^1(q^1, R^1) \geq V^1(q^2, R^2) \\ & \quad \quad V^2(q^2, R^2) \geq V^2(q^1, R^1) \\ & \quad \quad V^1(q^1, R^1) \geq V^1(0, 0) \\ & \quad \quad V^2(q^2, R^2) \geq V^2(0, 0). \end{aligned}$$

Thus, incentive compatibility constraints are added to the individual rationality (IR) constraints which were in the full information problem. Type 1's IR constraint will always bind at the optimum. With the single crossing assumption, type 2's IC constraint also binds.¹⁰ Denote the solution to this problem as $(\bar{q}^1, \bar{R}^1, \bar{q}^2, \bar{R}^2)$ and the contracts as $\bar{\alpha} = (\bar{q}^1, \bar{R}^1)$ and $\bar{\beta} = (\bar{q}^2, \bar{R}^2)$. (See Figure 1.) The firm offers a distorted contract to type 1's (one in which the slope of 1's indifference curve exceeds marginal cost). The type 2 contract is undistorted (the slope of V^2 equals marginal cost). Note that, in contrast to the full-

¹⁰ In comparison with the monopoly problem, the optimal tax problem replaces the profit function with a social welfare function, deletes the IR constraints, and adds the government budget constraint. The government budget constraint prevents separating in the final stage of the game types who reported the same type initially, but we explore this possibility in the game with additional contracts in Section 4.

information solution, the relative numbers of the two types of consumers influences the $\bar{\alpha}$ and $\bar{\beta}$ contracts because there is a tradeoff between the distortion on type 1 and the profit extracted from type 2.¹¹

When the number of customers is finite and the monopolist knows the exact distribution of types, the monopolist may take advantage of this knowledge by using an analog of Piketty's generalized tax schedules. Initially, as in Piketty we assume that the monopolist can commit to these schedules. We model this as follows. The customers will simultaneously choose among three possible strategies, to not purchase, to purchase as a high type, and to purchase as a low type. The monopolist will observe the number of customers who announce each type, denoted as (N^α, N^β) , where $N^\alpha + N^\beta \leq T$. The monopolist has announced in advance and is committed to a menu of contracts $C(N^\alpha, N^\beta)$ which state what a customer announcing low and one announcing high would get conditional on the number who announce each type. Customers are committed to purchase if they choose L or H regardless of what the other individuals revealed. The monopolist is restricted to selecting a contract menu which induces the resulting game among customers to be dominance solvable. It turns out best for the monopolist to induce dominance solvability in the following way. No matter how many other customers choose each of the three alternatives, it is better for a low-demand type to choose the strategy L. Given that all the low types reveal L, it must then be better for a high type to reveal H, no matter what the other high types choose. Note that truth telling for the high-demand types is not a dominant strategy, although it is for the low-demand types.

¹¹ The monopolist may only wish to sell to a single type in this problem, in which case $(q^1, R^1) = (0, 0)$. This is effectively a corner solution to the Π^B problem.

Proposition 1: With a finite number of individuals, if full commitment is possible, the monopolist can achieve almost the same profits as under full information by specifying a menu of contracts which induces a dominance-solvable game among consumers.

Proof: Let $\alpha(\varepsilon)$ and $\beta(\varepsilon)$ be defined as undistorted contracts for types 1 and 2 (the slope of V^1 equals marginal cost at (q^1, R^1)) such that $V^1(\alpha(\varepsilon)) = V^1(\hat{\alpha}) + \varepsilon$ and $V^2(\beta(\varepsilon)) = V^2(\hat{\beta}) + \varepsilon$, for small ε , where $\hat{\alpha}$ and $\hat{\beta}$ are the full information contracts. Consider the menu of contracts defined as follows:

$$\begin{aligned} \alpha(N^\alpha, N^\beta) &= \alpha(\varepsilon), & \text{if } N^\alpha \leq N \\ &= (\gamma, R^\gamma), & \text{if } N^\alpha > N \\ \beta(N^\alpha, N^\beta) &= \beta(\varepsilon), & \text{all } N^\alpha, N^\beta \end{aligned}$$

where positive γ is chosen small enough so $V^2(\beta(\varepsilon)) > V^2(\gamma, R^\gamma)$.

Since $V^1(\hat{\beta}) < V^1(0,0)$, then $V^1(\beta(\varepsilon)) < V^1(0,0)$ will also hold for small ε , so that $V^1(\alpha(\varepsilon)) > V^1(0,0) > V^1(\beta(\varepsilon))$. Choose (γ, R^γ) so that $V^1(\gamma, R^\gamma) > V^1(0,0)$.¹² For a type 1 to announce truthfully therefore strictly dominates not purchasing or announcing as a high-demand type. If all type 1's announce truthfully, then a type 2 who is truthful gets $\beta(\varepsilon)$. A type 2 who does not purchase gets $(0,0)$ and one who claims to be a type 1 gets (γ, R^γ) . Since $V^2(\beta(\varepsilon)) > V^2(\gamma, R^\gamma) > V^2(0,0)$, it is then a strictly dominant strategy for 2's to be truthful if all 1's are truthful. At the unique equilibrium where everyone is truthful, the profits can be made arbitrarily close to those at full information by setting ε , γ , and R^γ arbitrarily close to 0. **QED**

¹² The (γ, R^γ) contract lies between the V^1 indifference curve through $(0, 0)$ and the ray through $(0, 0)$ with a slope equal to marginal cost. Thus, this contract is profitable for the monopolist, guaranteeing that the monopolist earns nonnegative profit in all states of the world.

This solution effectively replicates Piketty's result for generalized tax schedules. The difficulty with this mechanism lies in what happens off the equilibrium path. The monopolist has committed to sell essentially nothing to type 1 consumers when more than N consumers (and therefore, at least some type 2's) claim to be type 1's. The monopolist earns zero profit from each of these consumers when that occurs. If it cannot commit to this outcome, then it will offer another contract to these individuals, possibly the original $\hat{\alpha}$ contract. More generally, if the monopolist cannot commit, then any contract which does not maximize profit given the monopolist's beliefs about the distribution of consumer types in any information set would be subject to revision. We consider the possibility of revision in the next sections.

3. Sustaining the Full-Information Outcome under Partial Commitment

Both the standard asymmetric information optimum and the solutions where contracts depend on how many customers select each type require the monopolist to be able to commit to the contracts announced. If there is no possibility of commitment, no information revelation by customers is possible. Consider any situation in which separation occurs. The monopolist, now knowing the type of each customer, would adjust the contracts to ones which extract all surplus from each type. Knowing that this will occur, all consumers would claim to be the lowest demand type. In the standard case, the customers would then get the contract promised to the low type. In the Piketty case, the monopolist has threatened to sell nothing but will not carry out the threat without commitment. Furthermore, at the start, the monopolist will recognize that everyone will reveal the same type and thus that it will be in effect limited to offering a single contract to everyone. The most profitable single contracts are either just $\hat{\alpha}$ which will be accepted by everyone or just $\hat{\beta}$ which will induce purchases by

only the high-demand types. Either of these yields less profits than the Piketty-type solution and no more than the conventional second-best solution.

Even if the monopolist cannot fully commit to a menu of contracts, it may not be limited to the no-commitment outcome. Legally binding contracts can be altered if this benefits the customer since there would be no one with standing to object. Revisions that would make a customer worse off would yield a legal challenge. On the equilibrium path (after truthful revelation), all customers who accepted a particular contract are identical so replacing an offer with a better one is straightforward. Customers of the other type might now desire the new offer but would have no legal recourse. Off the equilibrium path (after misrevelation by some customers), different types of customers have revealed themselves to be of the same type. The monopolist can offer a new contract that makes those who revealed truthfully better off even if it harms those who misrevealed. Customers who misrevealed are able to prove harm from the replacement contract only by admitting to their misrevelation, and thus losing any legal standing to sue.¹³

We refer to this commitment by the monopolist as partial commitment because it only guarantees consumers a lower bound on their payoff (and only if they reveal truthfully). We now consider the game with partial commitment in which we impose this as a requirement both on and off the equilibrium path.

The structure of this game is as follows. First, the monopolist (as before) announces the menus of contracts, $C(N^\alpha, N^\beta)$. Second, consumers simultaneously choose either the α or β

¹³ The reporting of types by consumers in our mechanism differs from its use in the revelation principle. Under the revelation principle, one constructs a direct mechanism game as a formal description of an indirect mechanism without loss of generality. In our mechanism, switching from reports of types to actions by consumers would in general result in a different game. We explore this more fully in Hamilton and Slutsky [2003].

menus or forgo the opportunity to buy from the monopolist. In the latter case, there is no possibility of the monopolist later attracting such customers by offering a revised contract (we consider alternatives later). After publicly revealing N^α and N^β (the number of consumers who choose the α and β menus), the monopolist can offer substitute contracts to one or both types. The partial commitment on the part of the monopolist is a commitment that the replacement contracts will satisfy:

$$V^1(\alpha'(N^\alpha, N^\beta)) \geq V^1(\alpha(N^\alpha, N^\beta)) \quad \text{and} \quad V^2(\beta'(N^\alpha, N^\beta)) \geq V^2(\beta(N^\alpha, N^\beta)).$$

The revised contract must not harm consumers of the type for whom that contract is intended, but consumers who misreported their types may be made worse off as a result of the contract revision. Because of the legal nature of the contract, the choice of a menu will be binding on a consumer—he cannot claim to be a different type after the contract revision. In any case, no one could gain by doing that in equilibrium. The customer can even be compelled not to drop out. After the monopolist chooses these new contracts, all N^α consumers who chose the α menu receive $\alpha'(N^\alpha, N^\beta)$, and all N^β consumers who chose the β menu receive $\beta'(N^\alpha, N^\beta)$.

Note that the game has $(T+1)(T+2)/2$ information sets, where each information set corresponds to an allowable value of (N^α, N^β) .¹⁴ The monopolist knows how many consumers claim to be of each type, but in most information sets it cannot identify precisely the number who claim to be of each type who are reporting truthfully. For example, if $N^\alpha = N + 1$, all type 1's and one type 2 may have claimed to be type 1, or $N-1$ type 1's and two type 2's may have claimed to be type 1, and so forth. When the equilibrium of the game will be truthful revelation, the only information set reached with positive probability is that with $N^\alpha = N$ and

¹⁴ Because of the possibility of opting out, this number of information sets is greater than it would be if all consumers must purchase some option..

$N^\beta = T - N$. The limiting beliefs in all other information sets correspond to assuming the minimum misrevelation necessary (in absolute numbers). This means that any information set is believed to be reached by having the fewest number of people stating they are of the other type. For example, when $N^\alpha = N + 1$ and $N^\alpha + N^\beta < T$, the monopolist believes that a single type 2 has misrevealed and all nonparticipants are type 2's. Using beliefs formed in this way satisfies the properties of a sequential equilibrium (Kreps and Wilson, 1982).

Our first observation is that the revised contracts will always be undistorted contracts.

Lemma 1: For every N^α and N^β , a menu of revised contracts $\alpha'(N^\alpha, N^\beta)$ and $\beta'(N^\alpha, N^\beta)$ is consistent with the requirements of partial commitment if and only if the contracts have zero distortion for their intended type (the slope of V^1 at $\alpha'(N^\alpha, N^\beta)$ and the slope of V^2 at $\beta'(N^\alpha, N^\beta)$ equal marginal cost).

Proof: After consumers have chosen their menus, there are no further self-selection constraints. Thus, under partial commitment the $\alpha'(N^\alpha, N^\beta)$ contract is the solution to:

$$\text{Max}_{q^1, R^1} \quad N^\alpha [R^1 - cq^1] \quad \text{s.t.} \quad V^1(q^1, R^1) \geq V^1(\alpha(N^\alpha, N^\beta)),$$

while the $\beta'(N^\alpha, N^\beta)$ contract is the solution to:

$$\text{Max}_{q^2, R^2} \quad N^\beta [R^2 - cq^2] \quad \text{s.t.} \quad V^2(q^2, R^2) \geq V^2(\beta(N^\alpha, N^\beta)).$$

It is immediate that $\partial R^i / \partial q^i = c$ along V^i at the solution to these problems. Hence, the contracts must be undistorted.

Conversely, for any N^α and N^β , take any pair of undistorted contracts α' and β' . No revisions are possible. The isoprofit contour is tangent to each types indifference curve at the

bundle intended for that type. Any deviation which does not lower profits must lower the welfare of some customer who has revealed truthfully and hence will be blocked. **QED**

Given this, the contracts announced at the earlier stage can be viewed as guarantees of utility levels to consumers since everyone will know that, whatever specific (q,R) pair is offered for any given N^α and N^β , what will be implemented will be the nondistorted contracts yielding the same utility levels. Denote these utility levels as $V^1(N^\alpha, N^\beta)$ and $V^2(N^\alpha, N^\beta)$ for all allowable N^α and N^β . Since there is only one undistorted bundle on each indifference curve from the assumption of strict quasiconcavity, there is a one-to-one relationship between each type's set of undistorted bundles and its utility levels. Thus, we can translate the incentive constraints into limits on the utilities obtained at different information sets. For example, consider the information set (N^α, N^β) and assume the monopolist desires to give type 1's utility equal to $V^1(N^\alpha, N^\beta)$. If type 2's obtain too high a utility level at the information set $(N^{\alpha-1}, N^{\beta+1})$, then type 1's would have an incentive to lie. Thus, for each (N^α, N^β) , there is an upper bound on $V^2(N^{\alpha-1}, N^{\beta+1})$ to satisfy incentive compatibility. The higher is $V^1(N^\alpha, N^\beta)$, then the greater is the upper bound on $V^2(N^{\alpha-1}, N^{\beta+1})$. Hence we can write the incentive constraint as $V^2(N^{\alpha-1}, N^{\beta+1}) \leq F(V^1(N^\alpha, N^\beta))$ for some increasing function F which depends on preferences. Similarly, the other incentive constraint is $V^1(N^{\alpha+1}, N^{\beta-1}) \leq G(V^2(N^\alpha, N^\beta))$ for some increasing function G .

We can now find the sequential equilibrium of the subgame in which all consumers report their types truthfully.

Proposition 2: The monopolist can specify contracts which satisfy partial commitment, induce truth telling as the dominance solvable equilibrium, and earn approximately the same profits as under full information.

Proof: Consider the following menu of contracts offered by the monopolist. Anybody who reveals themselves as a high, no matter what the pattern of revelation by others, receives the contract $\beta(\epsilon)$, that which is a undistorted bundle for a high and which yields ϵ in utility to a true high, $V^2(\beta(\epsilon)) = \epsilon$. Note that a true low who reveals high gets negative utility since $V^1(\beta(\epsilon)) = -K < 0$. The assignment to lows depends upon the number of highs. In any information set in which the number who reveal high is at least the true number ($N^{\beta} \geq T-N$), the bundle given to a low is $\alpha(\epsilon)$, the nondistorted bundle for a low such that $V^1(\alpha(\epsilon)) = \epsilon$. In these circumstances, a true high who reveals low would get a positive utility, $V^2(\alpha(\epsilon)) = M$. In any information set in which the number who reveal as high is too few ($N^{\beta} < T-N$), anyone who says low is given $\bar{\alpha}$, an undistorted bundle for a low which yields negative utility to both a true high and a true low, $V^1(\bar{\alpha}) = -X$ and $V^2(\bar{\alpha}) = -Y$. The payoffs facing true lows and highs are shown in Table 1.

For any low, in any circumstance, not participating or claiming to be low strictly dominates playing H. Hence, no low will ever play high. Then, for any true high, $\hat{N}^{\beta} \leq T-N-1$ must hold where \hat{N}^{β} is the number of others who reveal themselves as high. In this case, revealing truthfully as a high strictly dominates the strategies of revealing as a low or opting out. Hence, all true highs reveal truthfully. Therefore, after two rounds of eliminating dominated strategies, any low perceives that $\hat{N}^{\beta} = T-N$ so that revealing truthfully is better than not participating. From Lemma 1, these contracts satisfy partial

commitment. The equilibrium of the game is truth telling which results in the $(\alpha(\epsilon), \beta(\epsilon))$ contracts which yield approximately the full information profits to the monopolist. **QED**

Note that while the outcome of full surplus extraction is the same as in Proposition 1, the possibility of limited reoptimization complicates the reasoning that customers must carry out. Truthful revelation is no longer a dominant strategy for low demanders. High demanders do not have a dominant strategy either, until we rule out the possibility of low demanders claiming to be high. Thus, full advantage is taken of requiring only dominance solvability instead of having truth as a dominant strategy. Furthermore, we cannot implement our outcome as a dominant strategy mechanism because low types will only prefer to take the contract for lows instead of opting out when all high types take the contract for high types. Since high types gain by taking the low contract when some low types take the high contract, it can never be a dominant strategy for high types to take the high contract.

In the Appendix, we establish that Proposition 2 extends to the case of more than two types of consumers. In the remainder of the paper, we consider only the two-type case to keep the notation understandable.

In the optimal income tax problem, Hamilton and Slutsky [2003] show that partial commitment is not enough to allow the social planner to sustain the full information solution. Why can a monopolist succeed in doing this? Two factors are significant in the difference. First is the absence of a budget constraint for the monopolist. This gives the monopolist more flexibility than the social planner in specifying off equilibrium path contracts. And second is the difference in the preferences of the monopoly and the social planner. The social planner tends to have at least partially similar preferences to the individuals in the society. Once types have been revealed, this means that there may exist significant agreement in what changes can

be made to initial proposals consistent with partial commitment. The monopolist, however, has essentially opposing preferences to customers in that, at an efficient point, gains in profits require reductions in utility. Thus, the “penalty” contracts are credible, while they would not be for a social planner in the optimal tax problem. This limits the modifications that can be made off the equilibrium path and helps sustain the full information outcome in equilibrium.

4. Offers to Consumers Who Reject Both Menus Initially

In the game studied in Proposition 2, consumers who opted out gave up any opportunity to buy from the monopolist. Given that we study partial commitment because consumers and the firm cannot commit not to revise contracts, it is appropriate to study what happens when the monopolist can make offers to consumers who declined either menu option initially.

We consider the case where the monopolist can only make a single take-it-or-leave-it offer to consumers who opted out in the first stage of the game. The game is identical to that studied in Proposition 2 except for the information sets reached after some consumers opt out. For each (N^α, N^β) pair such that $N^\alpha + N^\beta < T$, the monopolist will have a belief about the distribution of types among consumers who did not sign a contract at the final stage. Let μ denote the fraction of consumers who did not sign contracts that the monopolist believes to be low-demand types. The following lemma establishes that there are only two possible contracts offered to those who did not choose the α or β menus in the first stage.

Lemma 2: There exists a μ^* such that for $\mu \leq \mu^*$, the monopolist offers the $\hat{\beta}$ contract and only high-demand types choose to buy at the second opportunity, while for $\mu > \mu^*$, the monopolist offers the $\hat{\alpha}$ contract and both types choose to buy at the second opportunity.

Proof: Since these consumers opted out at the first stage, there are no partial commitment restrictions on these contracts. Let (q^0, R^0) be the contract offered to those who initially opted out. If $V^1(q^0, R^0) > 0$, all consumers will buy at this stage, while if $V^2(q^0, R^0) > 0 > V^1(q^0, R^0)$, only high-demand types will buy at this stage. There are thus two local maxima to the profit maximization problem, and they are the $\hat{\alpha}$ and $\hat{\beta}$ contracts. Let $\Pi(\hat{\alpha})$ and $\Pi(\hat{\beta})$ denote per-customer profit from the $\hat{\alpha}$ and $\hat{\beta}$ contracts. Hence, the monopolists' profits from these customers are $(T - N^\alpha - N^\beta) \Pi(\hat{\alpha})$ or $(1 - \mu)(T - N^\alpha - N^\beta) \Pi(\hat{\beta})$. If $\mu < \mu^* = \frac{\Pi(\hat{\beta}) - \Pi(\hat{\alpha})}{\Pi(\hat{\beta})}$, the $\hat{\beta}$ contract yields greater profit from customers who opted out. **QED**

As in other cases where only some consumers deviate from truthful revelation, the beliefs to consider for sequential equilibrium are those where the minimum number misrevealed. Thus $\mu = \min \left\{ \max \left\{ 0, \frac{N - N^\alpha}{T - N - N^\beta} \right\}, 1 \right\}$ for $T > N^\alpha + N^\beta$ is the monopolist's belief about consumers who opt out. When all consumers opt out initially, $\mu = \frac{N}{T - N}$ is the only consistent belief. In constructing a mechanism relying on iterated dominance, whether μ^* is greater or less than $\frac{N}{N + 1}$ will be crucial. If all low types and one high type opt out in the first stage, $\mu = \frac{N}{N + 1}$. If $\frac{N}{N + 1} < \mu^*$, then the monopolist offers the $\hat{\beta}$ contract to those who opted out and only high types buy at the second opportunity. If $\frac{N}{N + 1} > \mu^*$ and all low types

opt out initially, then a high type who opts out initially would obtain the \hat{a} contract at the second opportunity.¹⁵

Proposition 3: If the monopolist can make a single offer to all consumers who opted out when offered the choice between two contracts menus, there exists a dominance – solvable mechanism with truth telling which earns approximately the same profits as under full information.

Proof: See Table 2 for payoffs to both types of consumers. As before, it is a dominated strategy for low types to claim to be high types. The payoffs listed in all remaining contingencies for high types declaring high makes that a dominant strategy. Then $N^a > N$ can never occur, so claiming to be low dominates opting out for the low types. Since both types have to do better when telling the truth than being identified as the only consumer to opt out, $V^1 = V^2 = 2\varepsilon$ in the truth-telling equilibrium. **QED**

Thus, whether or not the monopolist can make new offers to those who opted out does not have a major influence on the equilibrium outcome.

5. Limitations on Contracts

A critical feature of Propositions 2 and 3 is the ability of the monopolist to impose a large penalty when not enough consumers claim to be the high type. This is what makes it possible to construct contracts with little surplus and make it a dominant strategy for high

¹⁵ In this game, there are additional incentive constraints that $V^1(N, T-N)$ must be at least as great as 1's continuation payoff after initially opting out at the first stage; similarly, $V^2(N, T-N)$ must be at least as great as 2's continuation payoff after initially opting out at the first stage.

types to be truthful. In equilibrium, no consumer regrets not opting out, but off the equilibrium path, consumers may end up worse off than if they had never signed a contract.

If the monopolist is providing an essential commodity to consumers, it may well be subject to regulation. One form the regulation could take is that the monopolist's offer must not violate the individual rationality constraint for truthful consumers in any state of the world and not just in equilibrium.

We model our second partial commitment framework where the monopolist can make new offers to consumers who initially reject both of the contract menus. In all information sets, the options must satisfy the individual rationality constraint. This restriction on off-the-equilibrium path contracts affects how much the monopolist can extract from the high-demand consumers. Proposition 4 describes the best pair of contract menus that the monopolist can offer with this restriction.

Proposition 4: When all elements of a contract menu must satisfy the individual rationality constraint for truthful consumers, the monopolist earns less profit than in the full-information outcome, even though the contracts are undistorted. In the best dominance-solvable outcome for the monopolist, low-demand consumers receive an efficient contract which is ε better than $\hat{\alpha}$, and high-demand consumers receive an efficient contract which is δ better than $\tilde{\beta}$ where $\tilde{\beta}$ is the efficient contract for high types such that $V^2(\tilde{\beta}) = V^2(\hat{\alpha})$.

Proof: See Figure 2 and Table 3 for the payoffs to both types. In order to implement a dominance-solvable outcome, as before, offering the zero-surplus contract to anyone who claims to be high makes reporting high a dominated strategy for the low types. If too many consumers claim to be the low type, the monopolist is constrained to give the low types a

contract such that $V^1(q, R) \geq V^1(0, 0)$. Any such contract other than $\hat{\alpha}$ would be revised to $\hat{\alpha}$. Thus, high types can receive M (the utility high types receive from $\hat{\alpha}$). In order to make reporting low a dominated strategy for high types when no low types report high, the payoff to high in this case must be greater than M . Thus, offering a contract δ better than $\tilde{\beta}$ will induce truthful revelation by highs when no lows claim to be high.

For lows, opting out is dominated by reporting as low when all highs report truthfully as long as they get 2ε of surplus from saying low. A contract marginally better than $\hat{\alpha}$ maximizes the monopolist's profit under this constraint.

Thus, the equilibrium outcome is truthful revelation with contracts which are ε and δ better for their purchases than $\hat{\alpha}$ and $\tilde{\beta}$. The monopolist's profit is lower than in the earlier mechanisms because it cannot punish those who claim to be low when too many consumers make that claim. **QED**

Thus, restrictions on the contracts off the equilibrium path reduce the monopolist's profit relative to the full-information case. These restrictions also reduce the monopolist's profit below that in the conventional asymmetric information problem. We can see this from the fact that the incentive compatibility constraint for the high types is still satisfied, but the fact that the monopolist can only partially commit prevents it from distorting the contract for the low types, as it does to maximize profits in the conventional problem.

Our mechanism uses a particular order of deletion of dominated strategies, starting with low types never wanting to claim to be high. The monopolist cannot do better with a different order of deletion since a high type would always benefit by claiming to be low when some low type claimed to be high. Since a high type can always earn surplus equal to M by

claiming to be low given the contract restrictions, any mechanism which induces high types to be truthful must give them at least this amount of surplus. It would be difficult to guarantee uniqueness of the Nash equilibrium for any mechanisms which were not dominance solvable. While other mechanisms might earn greater profit for the monopolist, they are likely to have multiple equilibria.

One reason to consider requiring the monopolist to satisfy the individual rationality constraints is the problem with a regulated monopolist selling an essential good. In such cases, it may not be possible for consumers to opt out, since they need to buy some of the good to survive. Essentially the same outcome as above would still arise. In order to prevent highs claiming to be the low type, the monopolist must give truthful highs more surplus than they would get with the $\hat{\alpha}$ contract (the worst contract the lows can get in equilibrium).

6. Some Extensions

One obvious weakness of our approach is the reliance of the monopolist's precise knowledge of the distribution of consumer types. While this precise knowledge is necessary to extract all consumer surplus, conditioning contracts on consumer reports can benefit the monopolist in other circumstances. Suppose that the monopolist knew that N_1 consumers have a type lying in an interval from θ_0 to θ_1 and that N_2 consumers have a type lying in an interval from θ_2 to θ_3 , where $\theta_0 < \theta_1 < \theta_2 < \theta_3$ and higher values of θ correspond to shifts out in the demand curve of the consumer. Within each interval, the monopolist only knows the distribution from which types in that interval are drawn. The conventional solution would apply a single self-selection constraint to the problem and the θ_2 consumer would be indifferent between the contract for θ_2 and θ_1 .

While the monopolist cannot observe misrevelations within the two intervals, the monopolist can observe misrevelation when a consumer claims to be of a type in the other interval. Consider the following mechanism. Offer contracts with zero surplus for types θ_0 and θ_2 if N_1 consumers claim to be in the low interval and N_2 consumers claim to be in the high interval, and with negative surplus for type θ_0 if more than N_1 consumers claim to be in the low interval. For simplicity, assume that all consumers reporting to be in the low interval receive the same contract as the θ_0 type and all consumers reporting to be in the high interval receive the same contract as the θ_2 type.¹⁶ In comparison with the conventional solution, this mechanism extracts less profit from consumers in $(\theta_0, \theta_1]$ because they get all their information rent relative to the θ_0 type. It also extracts less additional profit from consumers in $(\theta_2, \theta_3]$ relative to the θ_2 type. The gain comes from extracting full surplus from the θ_2 type and, as a consequence, reducing the information rent of at least some consumer with types greater than θ_2 .¹⁷ If the length of the intervals were small relative to the distance between them, conditioning offers on the number claiming to be in the low interval could be more profitable to the monopolist. Thus, in parallel with Levine and Pesendorfer (their Theorem 2, p. 1164-5), some level of noise need not eliminate a solution close to the full surplus extraction outcome, even though our solution does not require the same type of precommitment.

Even with a finite number of types, the addition of noise to the problem does not necessarily prevent the monopolist from doing almost as well as when it knows the precise

¹⁶ The mechanism could offer a menu within intervals as in the conventional second-best solution, but this would complicate the exposition unnecessarily.

¹⁷ Using a schedule of contracts within the high interval would permit the monopolist to extract more profit from all the high types.

number of each type. Consider the model of Section 3 where consumers who opt out cannot be offered contracts in the future. Suppose that there is a small probability ω that one high type selects the low contract regardless of the payoffs (this player simply behaves irrationally). It is still the case that no low types will choose the high type contract (because opting out has a higher payoff in all states), and therefore, after deleting this possibility, it is a dominant strategy for high types to choose the high contract. Let $\tilde{\varepsilon}$ be the payoff to low types taking the low contract when everyone chooses the contract for their type. For lows to prefer the low contract to opting out requires that $(1 - \omega)\tilde{\varepsilon} + \omega(-X) > 0$. For small ω , the monopolist needs only to give all consumers who choose the low contract a small surplus. Thus, the monopolist's profit is close to the full-extraction level. Similarly, a monopolist who believed that, with small probability, there were $N+1$ low types (instead of N) could use the same mechanism to induce truthful revelation and extract almost all surplus.

We can relax the precision on the distribution of types further. Let π_k equal the probability that the number of low types equals k . Our theorems in Sections 3-5 consider the case of $\pi_N = 1$. Suppose instead that $\pi_N = 1 - \gamma^L - \gamma^H$, $\pi_{N-1} = \pi_{N+1} = \omega$, $\pi_{N-2} = \pi_{N+2} = \omega^2$, and $\pi_{N-k} = \pi_{N+k} = \omega^k$ with $\gamma^L = \sum_{k=1}^{N-1} \omega^k$ and $\gamma^H = \sum_{k=1}^{T-N-1} \omega^k$. Let the contract for those choosing high and the contract for those choosing low when $N^\beta > N$ be the same as in Section 3. Again, no lows will ever wish to claim to be high. For lows to prefer to choose low to opting out requires that they receive a small surplus when N consumers claim to be low to compensate for the punishment $(-X)$ when more than N choose the low contract. For small values of ω , the monopolist again extracts almost the entire full-information surplus.

7. Concluding Remarks

We have studied the effects of allowing a monopolist to condition nonlinear pricing contracts on the number of consumers who choose each of the menus of contracts. In addition, we require that all contract options, on and off the equilibrium path, be robust to revisions which benefit both the monopolist and truthful consumers. The monopolist can extract all consumer surplus when there are no restrictions on the contracts which it can offer on and off the equilibrium path, even when it can only make partial commitments. In effect, the monopolist is able to implement its most preferred outcome from the set of undistorted outcomes. This contrasts with Hamilton and Slutsky's [2003] result for the optimal income tax problem where the budget constraint prevents the social planner from implementing some undistorted outcomes. When the monopolist can only offer contract menus which do not violate the individual rationality constraints on and off the equilibrium path, the dominance-solvable outcome has undistorted contracts, but these contracts are less profitable for the monopolist than the conventional distorted ones.

Unlike the simplest mechanisms where the monopolist exploits his knowledge of the precise distribution of consumer types, our mechanism does not have a multiplicity of Nash equilibria. Nor unlike mechanisms analogous to Piketty's generalized tax schedules do our mechanisms require the monopolist to commit to inferior outcomes off the equilibrium path. However, our mechanisms do require that the monopolist can offer contracts with negative surplus to truthful consumers off the equilibrium path. When we require that truthful consumers are never made worse off by participating in the mechanism, we find that the monopolist can only implement less profitable outcomes, albeit ones which are efficient.

The Bagnoli-Salant-Swierzbinski mechanisms [1989, 1995] also extract all or almost all consumer surplus by changing offers over time after some consumers accept the early high-price offers. While these solutions are subgame-perfect Nash equilibria, they are not unique. In particular, in the infinite horizon game, if some consumers with high valuations only buy when price falls near marginal cost, the equilibrium price will drop to marginal cost quickly and the outcome is closer to Coase's prediction. In contrast, our solutions are unique because of the dominance solvability property.

Other mechanisms (Crémer and McLean [1985, 1988], Spiegel and Wilkie [1999]) allow a monopolist to avoid paying information rents to consumers by exploiting correlation between consumer types. There are several differences between these environments and ours, as well as in the results. First, since everyone knows the precise number of each consumer type in our model, the correlation is quite special (perfectly negative correlation when there is one consumer of each type, for example). This special correlation is allowed by Crémer-McLean but is not required. Second, the Crémer-McLean mechanism requires consumer risk neutrality, while our mechanism does not impose this. Since our mechanism does work with risk-neutral consumers, it would appear that the contract revision is the crucial feature that prevents implementation by a dominant strategy mechanism. Third, Crémer-McLean find dominant strategy mechanisms in some cases and Bayesian Nash mechanisms in others, while we construct dominance-solvable mechanisms (which are intermediate in restrictiveness). Fourth, Spiegel and Wilkie decompose the payment portion of their mechanism into two parts—one which depends on one's own reported type and one which depends on the vector of others' reports. Such a decomposition is not possible in our mechanism since the surplus

received depends on the complete vector of reports. Thus, there appears to be no nesting relationship of our model with theirs.

One problematic property of the Crémer-McLean mechanism is that the transfers paid by agents may need to be quite large. In our model where the aggregate distribution of types is common knowledge, the vectors of conditional probabilities of one consumer's type given another's type are sufficiently different that that does not occur. Thus, while our mechanism does not incorporate very large payments by consumers in any state of the world, a Crémer-McLean mechanism would not require very large payments either in the circumstances we consider.

Appendix

The model in the text treats only the case of two types of consumers. Here, we show that the result in Proposition 2 generalizes to the case of a finite number of consumer types when the true number of consumers of each type is common knowledge. Our approach is inductive. Formally, we show here that, given the mechanism for n types ($n \geq 2$), we can construct a method for $n + 1$ types.

Consider three types of consumers with payoff function $V^1(q, R)$, $V^2(q, R)$, and $V^3(q, R)$, where type 3 indifference curves are steeper than type 2 indifference curves and type 2 indifference curves are steeper than type 1 indifference curves (as before) at any point of intersection. Since the mechanism worked for any pair of initial types whose preferences satisfied single crossing, this ordering is without loss of generality. Let N^1 , N^2 , and N^3 be the true number of each type.

Denote the contracts as α (intended for type 1), β (intended for type 2), and γ (intended for type 3). For all report vectors $(N^\alpha, N^\beta, N^\gamma)$, the γ contract is the undistorted contract yielding ε surplus to type 3 consumers. This contract thus provides negative surplus to type 1 and 2 consumers. If $N^\gamma \neq N^3$, the α and β contracts are the undistorted contracts for types 1 and 2 that yield $-\varepsilon$ surplus to type 3 consumers (and negative surplus to types 1 and 2). Since type 1 and 2 consumers can guarantee themselves zero surplus by opting out, choosing the γ contract is a dominated strategy for type 1 and type 2. Thus, type 3 consumers cannot gain by choosing the α and β contracts or opting out. If type 3 consumers all reveal their types truthfully ($N^\gamma = N^3$), the α and β contract menus in Proposition 2 induce truthful revelation as a dominance solvable equilibrium. Thus, the game with three types has a dominance solvable equilibrium with almost complete surplus extraction.

Similarly, for any $n+1$ consumer types, we can construct a dominance solvable equilibrium with almost complete surplus extraction in which reporting to be the highest demand type is a dominated strategy for all but the highest demand type. After deleting strategies involving misrevelation by lower types, the highest types strictly prefer to reveal truthfully. Then, the mechanism for n types induces truthful revelation by the remaining n types.

References

- Abreu, D., and H. Matsushima, 1992, Virtual Implementation in Iteratively Undominated Strategies: Complete Information, *Econometrica* 60: 993-1008.
- Abreu, D., and H. Matsushima, 1992b, Virtual Implementation in Iteratively Undominated Strategies: Incomplete Information, mimeo, Princeton University.
- Abreu, D. and A. Sen, Virtual Implementation in Nash Equilibrium, *Econometrica* 59: 997-1021.
- Bagnoli, M., S. Salant, and J. Swierzbinski, 1989, Durable-Goods Monopoly with Discrete Demand, *Journal of Political Economy* 97: 1459-1478.
- Bagnoli, M., S. Salant, and J. Swierzbinski, 1995, Intertemporal Self-Selection with Multiple Buyers, *Economic Theory* 513-526.
- Baron, D. and D. Besanko, 1987, Commitment and Fairness in a Dynamic Regulatory Relationship, *Review of Economic Studies* 54: 413-436.
- Brusco, S., 1998, Unique Implementation of the Full Surplus Extraction Outcome in Auctions with Correlated Types, *Journal of Economic Theory* 80: 185-200.
- Crémer, J., and R. McLean, 1985, Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent, *Econometrica* 53: 345-361.
- Crémer, J., and R. McLean, 1988, Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions, *Econometrica* 56: 1247-1257.
- Duggan, J., Virtual Bayesian Implementation, *Econometrica* 65: 1175-1199.
- Farrell, J. and E. Maskin, 1989, Renegotiation in Repeated Games, *Games and Economic Behavior* 1: 327-360.
- Fudenberg, D. and J. Tirole, 1990, Moral Hazard and Renegotiation in Agency Contracts, *Econometrica* 58: 1279-1319.
- Fudenberg, D. and J. Tirole, 1991, *Game Theory*, Cambridge, MA: MIT Press.
- Hamilton, J., W. B. MacLeod, and J.-F. Thisse, 1991, Spatial Competition and the Core, *Quarterly Journal of Economics* 106: 925-937.
- Hamilton, J. and S. Slutsky, 2003, Optimal Income Taxation with a Finite Number of Consumers, mimeo, University of Florida.

- Hammond, P., 1979, Straightforward Incentive Compatibility in Large Economies, *Review of Economic Studies* 46: 263-282.
- Kreps, D., and R. Wilson, 1982, Sequential Equilibria, *Econometrica* 50: 863-894.
- Levine, D, and W. Pesendorfer, 1995, When Are Agents Negligible?, *American Economic Review* 85: 1160-1170.
- Maskin, E. and J. Riley, 1984, Monopoly with Incomplete Information, *RAND Journal of Economics* 15: 171-196.
- Mirrlees, J., 1971, An Exploration in the Theory of Optimal Income Taxation, *Review of Economic Studies* 38: 175-208.
- Mookerjee, D. and S. Reichelstein, 1992, Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules, *Journal of Economic Theory* 56: 378-399.
- Piketty, T., 1993, Implementation of First-Best Allocations via Generalized Tax Schedules, *Journal of Economic Theory* 61: 23-41.
- Spiegel, Y., and S. Wilkie, 2000, Optimal Multiproduct Nonlinear Pricing with Correlated Consumer Types, mimeo, California Institute of Technology.
- Stiglitz, J., 1977, Monopoly, Non-linear Pricing and Imperfect Information: The Insurance Market, *Review of Economic Studies* 44: 407-430.
- Stiglitz, J., 1982, Self-Selection and Pareto Efficient Taxation, *Journal of Public Economics* 17: 213-240.

Low-Demand Consumers

		$\hat{N}^\beta < T - N$	$\hat{N}^\beta \geq T - N$
Announcement	O	0	0
	L	-X	ε
	H	-K	-K

High-Demand Consumers

		$\hat{N}^\beta < T - N$	$\hat{N}^\beta \geq T - N$
Announcement	O	0	0
	L	-Y	M
	H	ε	ε

(\hat{N}^β is the number of other customers who reveal themselves as highs.)

TABLE 1

The Payoff Matrix to Consumers under Partial Commitment
(Proposition 2)

Low-Demand Consumers

	$\hat{N}^a < N-1$ $\hat{N}^b < T-N$ $\mu > \mu^*$	$\hat{N}^a < N-1$ $\hat{N}^b < T-N$ $\mu < \mu^*$	$\hat{N}^b \geq T-N$ $\hat{N}^a < N-1$	$\hat{N}^a > N$	$\hat{N}^a = N-1$ $\hat{N}^b = T-N$ (truth)
O	$\varepsilon(\hat{\alpha} + \varepsilon)$	0(reject $\hat{\beta} + \varepsilon$)	$\varepsilon(\hat{\alpha} + \varepsilon)$	0(reject $\hat{\beta} + \varepsilon$)	$\varepsilon(\hat{\alpha} + \varepsilon)$
L	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	- X	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$
H	$-K_1(\tilde{\beta} + \delta)$	$-K_1(\tilde{\beta} + \delta)$	$-K_0(\tilde{\beta})$	$-K_0(\hat{\beta})$	$-K_0(\hat{\beta} + \varepsilon)$

High-Demand Consumers

	$\hat{N}^a < N-1$ $\hat{N}^b < T-N-1$ $\mu > \mu^*$	$\hat{N}^a < N-1$ $\hat{N}^b < T-N-1$ $\mu < \mu^*$	$\hat{N}^b \geq T-N$ $\hat{N}^a < N-1$	$\hat{N}^a > N$	$\hat{N}^b = T-N-1$ $\hat{N}^a = N$ (truth)
O	$M(\hat{\alpha} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$	$M(\hat{\alpha} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$
L	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	- Y	- Y
H	$M + 2\varepsilon(\tilde{\beta} + \delta)$	$M + 2\varepsilon(\tilde{\beta} + \delta)$	$0(\hat{\beta})$	$2\varepsilon(\hat{\beta} + 2\varepsilon)$	$2\varepsilon(\hat{\beta} + 2\varepsilon)$

(\hat{N}^a is the number of other customers who reveal themselves as lows and
 \hat{N}^b is the number of other customers who reveal themselves as highs.)

TABLE 2

The Payoff Matrix to Consumers under Partial Commitment When the Monopolist Can Make
 New Offers to Those Who Initially Opt Out
 (Proposition 3)

Low-Demand Consumers

	$\hat{N}^\alpha < N-1$ $\hat{N}^\beta < T-N$ $\mu > \mu^*$	$\hat{N}^\alpha < N-1$ $\hat{N}^\beta < T-N$ $\mu < \mu^*$	$\hat{N}^\beta \geq T-N$ $\hat{N}^\alpha < N-1$	$\hat{N}^\alpha > N$	$\hat{N}^\alpha = N-1$ $\hat{N}^\beta = T-N$ (truth)
O	$\varepsilon(\hat{\alpha} + \varepsilon)$	$0(\text{reject } \hat{\beta} + \varepsilon)$	$\varepsilon(\hat{\alpha} + \varepsilon)$	$0(\text{reject } \hat{\beta} + \varepsilon)$	$\varepsilon(\hat{\alpha} + \varepsilon)$
L	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$	$\varepsilon(\hat{\alpha} + \varepsilon)$	$2\varepsilon(\hat{\alpha} + 2\varepsilon)$
H	$-K_1(\tilde{\beta} + \delta)$	$-K_1(\tilde{\beta} + \delta)$	$-K_0(\tilde{\beta})$	$-K_0 + \varepsilon(\hat{\beta} + \delta)$	$-K_0 + 2\varepsilon(\hat{\beta} + \delta^*)$

High-Demand Consumers

	$\hat{N}^\alpha < N-1$ $\hat{N}^\beta < T-N-1$ $\mu > \mu^*$	$\hat{N}^\alpha < N-1$ $\hat{N}^\beta < T-N-1$ $\mu < \mu^*$	$\hat{N}^\beta \geq T-N$ $\hat{N}^\alpha < N-1$	$\hat{N}^\alpha > N$	$\hat{N}^\beta = T-N-1$ $\hat{N}^\alpha = N$ (truth)
O	$M(\hat{\alpha} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$	$M(\hat{\alpha} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$	$\varepsilon(\hat{\beta} + \varepsilon)$
L	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	$M + \varepsilon(\hat{\alpha} + 2\varepsilon)$	$M + \varepsilon(\hat{\alpha} + \varepsilon)$	$M + 2\varepsilon(\hat{\alpha} + 2\varepsilon)$
H	$M + 2\varepsilon(\tilde{\beta} + \delta)$	$M + 2\varepsilon(\tilde{\beta} + \delta)$	$0(\hat{\beta})$	$M + 2\varepsilon(\hat{\beta} + \delta)$	$M + 3\varepsilon(\tilde{\beta} + \delta^*)$

(\hat{N}^α is the number of other customers who reveal themselves as lows and
 \hat{N}^β is the number of other customers who reveal themselves as highs.)

TABLE 3

The Payoff Matrix to Consumers under Partial Commitment When Offers Cannot Violate
Participation Constraints in Any State of the World
(Proposition 4)