

Discussion Paper No. 784

STATIONARY SEQUENTIAL EQUILIBRIA
IN BARGAINING WITH
TWO-SIDED INCOMPLETE INFORMATION

by

Lawrence M. Ausubel*

and

Raymond J. Deneckere*

June 1988

* Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60208.

Financial support for this paper was provided in part by National Science Foundation Grants IST-86-09129 and SES-86-19012 and by the Kellogg Graduate School of Management's Esmark/Beatrice Research Chair. We would like to thank Roger Myerson and Robert Wilson for helpful discussions. Any errors that might exist in the paper, of course, are solely our responsibility.

1. Introduction

The bilateral bargaining problem in which a seller and buyer negotiate the exchange of a good (or, more generally, two rational individuals attempt to arrange any mutually beneficial transaction) is fundamental to the study of economics. It is enigmatic whether two rational agents will reach agreement and, if so, how they will divide the surplus, even under the simplifying assumptions that all information is common knowledge and that all dealing occurs at a single moment in calendar time. Yet the bargaining problem becomes still more complex when we introduce two additional, essential features into the model: (1) the bargaining process is explicitly dynamic and (2) each party possesses private information.

In this paper, we analyze the seller-offer bargaining game with two-sided incomplete information. A seller repeatedly proposes prices to a buyer, who accepts or rejects each offer. Each trader's valuation for the single good is unknown to the other player, but the distribution functions for seller and buyer valuations are common knowledge and have common support. Bargaining continues until such time that an offer is accepted. Although this game is one of the most streamlined infinite-horizon bargaining games with private information on both sides, we will see in this paper that it provides a fairly rich environment for studying the nature of sequential bargaining.

One of the basic issues we explore is the relationship between static and dynamic bargaining. Research by other authors had warned that, in a dynamic context, players would necessarily suffer a substantial loss in ex-ante expected gains from trade (relative to the ex-ante efficient mechanism), due to their "inability to commit to the static mechanism"

(Cramton, 1988, Section 6). The intuition behind this critique is that optimal mechanism design requires that, sometimes, the buyer's valuation will strictly exceed the seller's and yet the players will trade with probability zero.¹ Such an arrangement--plausible in a static context--might conflict with our notion of sequential rationality in a dynamic game if it necessitates "that the bargainers are able to commit to walking away from the negotiating table, even when it is common knowledge that the gains from trade are positive" (Cramton, 1984, p. 591). Roughly the same argument is implicit in Fudenberg, Levine and Tirole (1985), and Chatterjee and Samuelson (1986, 1987).

Another fundamental issue we study is the role of stationarity in bargaining solutions. Many authors have argued that stationarity--the notion that history should matter only insofar as it is reflected in the current "state"²--is an interesting property to find in sequential equilibria of bargaining games.³ It had been hoped that a restriction to stationary sequential equilibria would provide attractive solutions to the infinite-horizon bargaining game with two-sided incomplete information, especially as the time interval between successive offers was allowed to converge to zero (so that the extensive form concealed arbitrarily little exogenous precommitment). Unfortunately, efforts to construct such equilibria were frustrated by what has become known as the Coase conjecture: various stationarity restrictions in bargaining games with one-sided incomplete information imply that the offers of the seller (the uninformed party) rapidly converge to the lowest buyer valuation.

In this paper, we establish a series of results concerning stationary sequential equilibrium in the seller-offer bargaining game. One major

finding is that there is considerable loss in generality and plausibility in examining equilibria which satisfy the additional restriction that the Coase conjecture holds for the lowest seller type. This additional restriction does delineate what might have seemed to be a natural class of sequential equilibria and, indeed, the equilibria in Cramton (1984) and Cho (1988) satisfy the latter requirement. Unfortunately, we demonstrate in Theorem 2 that, when the supports of the seller and buyer distribution functions have a common minimum value (and the seller distribution does not possess a mass point at the bottom of its support), the restriction has a rather severe implication: when the time interval between offers is made brief, the bargaining outcome approximates zero trade. That pessimistic conclusion might be thought to support the concern that inefficiency is inherent in sequentiality.

However, all is not lost in the study of stationary sequential equilibria of this game, provided we forego the above property. Instead, we consider equilibria in which the lowest seller type does not fully reveal her valuation before beginning to sell (and in fact never fully reveals), and so the Coase conjecture need not hold. Our challenge is then to identify a class of static mechanisms which are implementable by stationary sequential equilibria, in the sense that there exists a sequence of such equilibria whose outcomes converge (in measure) to the desired static mechanism. Since the seller makes all the offers, a natural definition to propose is the following: a seller-first mechanism is a static mechanism which would maintain its incentive compatibility even if the seller were required to publicly reveal her type before the buyer made his announcement.

Although the seller-first mechanisms are natural candidates for

outcomes of seller-offer bargaining games, we demonstrate that not all seller-first mechanisms are implementable. The reason for this is that the sequential nature of the game does place some restriction on the form of the equilibrium: certain deviations by the seller cannot be deterred by even the most adverse buyer inferences. The avoidance of unstoppable deviations motivates our definition of sequentially seller-first mechanisms. We are then able to establish the main result of the paper, Theorem 3: (essentially) every sequentially seller-first mechanism is implementable in stationary, sequential equilibria.

We next ask: What is the seller-first mechanism which maximizes gains from trade? In Theorem 4, we assume that the distribution function of seller valuations is convex and the distribution function of buyer valuations is concave. Surprisingly, we demonstrate that the "monopoly mechanism" (in which each seller type charges her own monopoly price against the buyer distribution) is efficient within the class of seller-first mechanisms. Since it also satisfies the sequentiality requirement, we know that it is implementable by stationary sequential equilibria.

While the monopoly mechanism typically does not maximize expected gains from trade within the class of all incentive compatible bargaining mechanisms,⁴ it frequently comes fairly close. In the example where seller and buyer valuations are each uniformly distributed on a common interval, the monopoly mechanism realizes 8/9 of the gains from trade achieved by the efficient static mechanism. (See, also, Chatterjee and Samuelson, 1983; Myerson and Satterthwaite, 1983; and Myerson, 1985.)

Another reasonably attractive mechanism which can be implemented in stationary strategies is the single-price mechanism: for arbitrary price π ,

the players trade with probability one whenever the seller's valuation is less than π and the buyer's valuation exceeds π , and they trade with probability zero otherwise. In the double-uniform example, the "competitive mechanism" (where π equals the midpoint of the support) also realizes 8/9 of the gains achieved by the efficient static mechanism.

We thus argue that there is no need to despair of dynamic bargaining processes with asymmetric information. Rather appealing outcomes of the bargaining process are consistent both with stationarity and sequential rationality.

2. Literature Survey

The seminal paper examining infinite-horizon, offer/counteroffer bargaining games was that of Rubinstein (1982). He found that the alternating-offer game with complete information possesses a unique subgame perfect equilibrium. Moreover, he showed that stationarity is implied by the solution concept in such complete information games.

Much of the research on bargaining with one-sided incomplete information has centered about Coase's (1972) conjecture that a seller's price would drop toward the competitive level "in the twinkling of an eye."⁵ It turns out that the validity of the Coase conjecture hinges on whether a certain version of stationarity is required to be satisfied, and this in turn depends on the distributional assumptions. In the case of "the gap" (i.e., the seller's valuation equals s , the buyer's valuation is distributed over $[\underline{b}, \bar{b}]$ where $\underline{b} > s$, and the buyer's distribution function is Lipschitz-continuous at \underline{b}), sequential equilibrium implies stationarity and hence the Coase conjecture in the seller-offer game (Fudenberg, Levine and Tirole,

1985, and Gul, Sonnenschein and Wilson, 1986). In the case of "no gap" (i.e., $b \leq s$), stationarity is not implied by sequentiality. If it is assumed, the Coase conjecture holds (Gul, Sonnenschein and Wilson, 1986); if it is not assumed, a folk theorem holds in the seller-offer game (Ausubel and Deneckere, 1986, 1988).

One earlier (Cramton, 1984) and one contemporaneous (Cho, 1988) paper have been written using the same (two-sided incomplete information) extensive form which we analyze in this article. Peter Cramton studied the seller-offer game where seller and buyer valuations are each privately known and uniformly distributed. Cramton's equilibrium path displays "separation-over-time." Low-valuation seller types make revealing offers early in the game, while high valuations pool by making unrealistic offers (yielding zero sales) until times which credibly signal their types. After the initial revealing offer, the seller follows the successive-skimming strategy of the Coase conjecture (Stokey, 1981, and Sobel and Takahashi, 1983), where price descends exponentially toward her reported valuation. In deciding when to make her first serious offer, the seller must balance two competing factors: late revelation permits high prices and hence more profitable sales; early revelation requires reduced price but also lowers the cost of delay. The seller thus delays agreement in order to signal strength. Since delay hurts a low-valuation seller more than a high-valuation seller, strategic delay serves as a credible (albeit costly) screening device.

In-Koo Cho studied the seller-offer game with more general distributional assumptions, and proved the existence of a sequential equilibrium with a different description. All seller types reveal their valuations in the initial period. However, higher types (through their

higher offers) induce lower probabilities of acceptance. The Cho equilibrium is stationary in an extremely strong sense: subsequent behavior only depends on the seller's beliefs about the buyer's valuation. (In other words, the buyer's beliefs about the seller's valuation are not part of the "state.") Hence, the seller must prove her valuation anew in each subsequent period. Whereas Cramton's seller encountered an incentive constraint only in periods up until the one in which she first revealed, Cho's seller faces an analogous⁶ incentive constraint in all subsequent periods as well.

Interestingly, despite the fact that Cho's post-revelation equilibrium paths differ from those of bargaining games with one-sided incomplete information, Cho argues that (just as in Cramton's equilibrium) the Coase conjecture is satisfied for the seller with lowest valuation. It should therefore be observed that when the lowest seller and lowest buyer valuations coincide, and as the time interval between offers approaches zero, both equilibria yield arbitrarily little utility to the lowest (zero) seller type. Thus, as far as efficiency and "delay" are concerned, the Coase conjecture becomes a two-edged sword. It requires extreme efficiency (and "no delay") for the seller with zero valuation. But, in order to prevent the zero seller from mimicking higher valuations, this may necessitate extreme inefficiency (and long delay) for the seller with positive valuation. Indeed, in Section 6, we demonstrate the "no-trade theorem": as the response time approaches zero, strategic delay becomes arbitrarily costly.

Some other papers have described sequential equilibria with different qualitative properties, but in games with different extensive forms.

Chatterjee and Samuelson (1987) were able to demonstrate the existence of a unique Nash equilibrium in a game where the seller repeatedly makes offers but where the set of prices, the set of seller types, and the set of buyers are each doubletons. Since Chatterjee and Samuelson worked in a game with different extensive form, we find it illuminating to, instead, describe an equilibrium with a similar flavor in the game that we are studying in this article.⁷ Suppose that the seller distribution has support on $[\underline{s}, \bar{s}]$ and the buyer distribution has support on $[\underline{b}, \bar{b}]$, where $\underline{s} < \underline{b} < \bar{s} < \bar{b}$. In every period, in equilibrium, the seller charges either the price \bar{s} or the price \underline{b} . Charging the price \underline{b} has the effect of immediately ending the game, as the buyer accepts the price \underline{b} with probability one. However, whether the seller first cuts her price to \underline{b} , or whether the buyer first accepts a price of \bar{s} , is determined as if the players were engaged in a war of attrition. Delay, as usual, becomes a credible signal of strength. However, what is most striking about this equilibrium is that, despite the fact that only the seller is permitted to make offers, the equilibrium is almost fully symmetric. In every period n , there exists a cutoff s_n such that the seller cuts her price to \underline{b} if and only if her valuation is less than s_n . There also exists a cutoff b_n such that the buyer accepts the high price of \bar{s} if and only if his valuation exceeds b_n . Since $\underline{s} \leq s_0 < s_1 < s_2 < \dots < \underline{b}$ and $\bar{b} > b_0 > b_1 > b_2 \dots > \bar{s}$, each player's survival in the war of attrition causes his rival to update his beliefs (and makes the latter more pessimistic).⁸

Finally, Cramton (1988) constructed an equilibrium in a continuous-time model in which the bargainers use delay to signal their strengths. The equilibrium has the attractive features that it is stationary and fully

symmetric--the game permits each player the endogenous option to make the first offer. However, the extensive form has the property that, after a trader has made an offer, she must wait until her rival replies before making her next offer. As such, the extensive form may conceal a certain amount of commitment power that is not present in the standard offer/counteroffer bargaining game; it is not completely clear how to implement an analogous equilibrium in stationary strategies in the standard model.

3. Sequential Games and Static Mechanisms

Consider a trading situation in which individual 1 (henceforth referred to as the seller) owns a single indivisible object which she would like to sell to individual 2 (referred to as the buyer). We assume that s , the seller's valuation for the object, and b , the buyer's valuation for the object, are private information. Thus, each trader knows his own valuation at the time of bargaining, but treats his opponent's valuation as a random variable. In this paper, we assume that these random variables are distributed independently, according to the (common knowledge) distribution functions F_1 and F_2 , respectively. We also assume that the supports of these distribution functions coincide, and that they form an interval. Without further loss of generality, we will consider the case where $\text{supp } F_i = [0,1]$, for $i = 1,2$. For future reference, let us also denote by μ_i the measures corresponding to F_i , so that $\mu_1([0,s]) = F_1(s)$ and $\mu_2([0,b]) = F_2(b)$. Define a bargaining mechanism to be a game in which each trader simultaneously reports his private information to a mediator, who then determines whether the good is transferred and how much the buyer must

pay the seller. A bargaining mechanism is characterized by two (Borel measurable) outcome functions: $p(s,b)$, which denotes the probability of trade given respective reports of s and b ; and $x(s,b)$, which denotes the expected transfer payment given those same reports. For any bargaining mechanism $\{p,x\}$, we define the following functions:

$$\begin{aligned}\bar{p}_1(s) &= \int_0^1 p(s,v_2)d\mu_2(v_2); & \bar{p}_2(b) &= \int_0^1 p(v_1,b)d\mu_1(v_1); \\ \bar{x}_1(s) &= \int_0^1 x(s,v_2)d\mu_2(v_2); & \bar{x}_2(b) &= \int_0^1 x(v_1,b)d\mu_1(v_1); \\ U_1(s',s) &= \bar{x}_1(s') - s\bar{p}_1(s'); & U_2(b',b) &= b\bar{p}_2(b') - \bar{x}_2(b').\end{aligned}$$

Thus, $\bar{x}_1(s)$ is the expected revenue and $\bar{p}_1(s)$ is the probability of losing the object, for a seller with valuation s . $U_1(s',s)$ denotes the expected utility to a seller of valuation s , when reporting a valuation s' to the mediator, and when the buyer reports truthfully. The quantities $\bar{p}_2(\cdot)$, $\bar{x}_2(\cdot)$, and $U_2(\cdot,\cdot)$ have a similar interpretation.

A bargaining mechanism is incentive compatible if honest reporting forms a Bayesian Nash equilibrium, i.e.,

$$(1) \quad \begin{aligned}U_1(s,s) &\geq U_1(s',s) & \forall s,s' \in [0,1], \\ U_2(b,b) &\geq U_2(b',b) & \forall b,b' \in [0,1].\end{aligned}$$

A bargaining mechanism is (interim) individually rational iff:

$$(2) \quad \begin{aligned}U_1(s,s) &\geq 0 & \forall s \in [0,1], \\ U_2(b,b) &\geq 0 & \forall b \in [0,1].\end{aligned}$$

Mechanisms satisfying both the individual rationality and the incentive constraints will be referred to as incentive compatible bargaining mechanisms (ICBMs).

In this paper, as in our previous work (Ausubel and Deneckere, 1988) with one-sided information, we study the relationship between incentive compatible bargaining mechanisms and sequential equilibria of infinite-horizon bargaining games. Specifically, the game form we analyze here has the seller make offers at discrete moments in time, spaced equally apart. Let z ($z > 0$) denote the time interval between successive offers. Then the seller makes bids at times $t = 0, z, 2z, 3z, \dots$. In each of these time periods, after hearing the seller offer, the buyer has an opportunity to accept or reject the offer. If the offer is accepted, the game ends and payoffs accrue. If the offer is rejected, the bargaining continues. Both players exhibit impatience, which is specified by a common discount rate r ($r > 0$). Hence, if the good is traded at time t for the price π between a seller with valuation s and a buyer with valuation b , the seller derives a surplus of $e^{-rt}(\pi - s)$ and the buyer obtains a surplus of $e^{-rt}(b - \pi)$.

Corresponding to any Nash equilibrium of the seller-offer game, there exists a sequential bargaining mechanism.⁹ Such a mechanism specifies a pair of outcome functions $t(\cdot, \cdot)$ and $x(\cdot, \cdot)$, where $t(s, b)$ denotes the time that the good will be transferred to the buyer, and $x(s, b)$ the expected payment to the seller, conditional on the reports (s, b) . Let $\bar{p}_1(s) = \int_0^1 e^{-rt(s, b)} d\mu_2(b)$ and $\bar{p}_2(b) = \int_0^1 e^{-rt(s, b)} d\mu_1(s)$. The incentive compatibility constraints for the sequential bargaining mechanism then are also (1). However, observe that acceptances are necessarily ex-post individually rational (along the equilibrium path), yielding the stronger

individual rationality constraints:

$$(3) \quad \begin{aligned} U_1(s,s) = \bar{x}_1(s) - s\bar{p}_1(s) &\geq 0, \quad \forall s \in [0,1], \\ be^{-rt(s,b)} - x(s,b) &\geq 0, \quad \forall s,b \in [0,1]. \end{aligned}$$

Thus, every Nash equilibrium of the seller-offer game induces, through the transformation $p(s,b) = e^{-rt(s,b)}$, a static bargaining mechanism which is incentive compatible.

Conversely, consider any ICBM $\{p,x\}$ which satisfies buyer ex-post individual rationality. Then, the mechanism $\{p,x\}$ suggests a time and a price at which (s,b) should trade, through the transformation:

$$(4) \quad \begin{aligned} t(s,b) &= -(1/r) \log p(s,b), \\ \pi(s,b) &= x(s,b)/p(s,b), \end{aligned}$$

where $\pi(s,b)$ is defined only when $p(s,b) > 0$. In one-sided incomplete information, where there is but a single seller type, transformation (4) parametrically defines a nonincreasing path of prices for the seller to charge over time. Moreover, every such price path is implementable in the seller-offer game (Ausubel and Deneckere, 1988).

However, in the case of two-sided incomplete information, matters are much more complicated. Typically, the times and prices generated by (4) cannot be interpreted as "price paths" for the seller. For fixed s , there will frequently exist b and b' such that $p(s,b) = p(s,b') > 0$ but $x(s,b) \neq x(s,b')$. Equation (4) then requires the same seller to charge two different prices at the same time, and for different buyer types to accept

these different prices. (The reader may wish to attempt this exercise for the Chatterjee-Samuelson (1983) mechanism, described by $\{p^5, x^5\}$ at the end of Section 8, below. This ICBM would ask $s \in [0, 3/4)$ to charge a continuum of prices, all at the same time.) Furthermore, even if $\{\pi(s,b), t(s,b)\}$ parametrically defines a price path for seller s , the path need not be monotone nonincreasing.

The interesting questions to be asked are: Under what conditions on $\{p,x\}$ does transformation (4) suggest declining seller price paths and when can close approximations to these price paths be sustained by sequential equilibria of the seller-offer game? These questions are nontrivial, for as will be shown in Section 4, there exist ICBMs which lend the price-path interpretation yet cannot be implemented. However, there also exists a nice class of ICBMs which are implementable.

4. Seller-First Mechanisms

We will now define a class of static mechanisms from which a menu of (nonincreasing) price paths can be read off quite naturally.

Definition 1: A seller-first mechanism is an ICBM $\{p,x\}$ satisfying:

- (i) $bp(s,b) - x(s,b) \geq bp(s,b') - x(s,b'), \quad \forall s,b,b' \in [0,1]:$
- (ii) $x(s,b) - sp(s,b) \geq 0, \quad \forall s,b \in [0,1];$ and
- (iii) $bp(s,b) - x(s,b) \geq 0, \quad \forall s,b \in [0,1].$

Condition (i) above requires that the buyer report truthfully, given knowledge of the seller's type. Thus, seller-first mechanisms are ICBMs for

which the buyer's incentive constraints are not upset if he were told which seller he was facing. In essence, this requires that a seller's first trade occurs either at a fully-revealing price, or that any seller types which have pooled together up until that point keep on pooling forever after. One way to visualize the bargaining situation is to imagine that the mediator, rather than soliciting simultaneous reports, first asks the seller to report her type within earshot of the buyer. Only then does the mediator require the buyer to reveal his valuation. Conditions (ii) and (iii) strengthen the requirements in the definition of an ICBM from interim IR to ex-post IR, for sellers and buyers, respectively. The reason for insisting on seller ex-post IR is explained in the following lemma.

Lemma 1: Let σ be a sequential equilibrium of the seller-offer game, and let $\{p,x\}$ be the induced static bargaining mechanism. Then if $\{p,x\}$ satisfies Definition 1(i) and (iii), $\{p,x\}$ also satisfies Definition 1(ii).

Proof: We will prove that (i) and (iii) imply that every seller type charges a (declining) price path, i.e., announces a (weakly) decreasing sequence of prices. This then immediately implies ex-post IR for the seller. Indeed, if the seller's price path ever dropped below her valuation, the seller could deviate by naming prices which induced zero sales, and would thereby avoid making unprofitable trade.

To see that the price paths are monotone, first observe that Definition 1(i) implies:

$$bp(s,b) - x(s,b) \geq bp(s,b') - x(s,b'), \text{ and}$$

(5)

$$b'p(s,b') - x(s,b') \geq b'p(s,b) - x(s,b).$$

Adding these inequalities yields: $(b' - b)p(s,b') \geq (b' - b)p(s,b)$. Thus, if $b' > b$, we have $p(s,b') \geq p(s,b)$.

Also, from (5) we see that:

$$b'[p(s,b') - p(s,b)] \geq x(s,b') - x(s,b) \geq b[p(s,b') - p(s,b)].$$

Letting $b' \rightarrow b$, this implies $dx_s(b) = bdp_s(b)$, for all $b \in [0,1]$, where $x_s(b) \equiv x(s,b)$ and $p_s(b) \equiv p(s,b)$. Since $x(s,0) = 0$, this implies $x(s,b) = \int_0^b rdp_s(r)$. Hence:

$$\partial/\partial b [x(s,b)/p(s,b)] = \{[bp(s,b) - x(s,b)]/p(s,b)\}^2 \{\partial/\partial b [p(s,b)]\},$$

whenever $p_s(b)$ is differentiable (and a similar expression holds at discontinuities). By Definition 1(iii), $bp(s,b) - x(s,b) \geq 0$, and hence $x(s,b)/p(s,b)$ is increasing in b . []

In the course of proving Lemma 1 we also established that, given $p(\cdot, \cdot)$, there is at most one function $x(\cdot, \cdot)$ so that $\{p, x\}$ is a seller-first mechanism. This function is given by $x(s,b) = \int_0^b rdp_s(r)$. Hence we obtain the following corollary to Lemma 1:

Lemma 2: Suppose $\{p^1, x^1\}$ and $\{p^2, x^2\}$ are seller-first mechanisms with

$p^1 = p^2$. Then $x^1 = x^2$.

It is worth noting that the set of all seller-first mechanisms has a simple structure: it is a nonempty, closed, convex subset of the set of all incentive compatible bargaining mechanisms. It is also a fairly rich set, as the next three examples demonstrate.

Example 1: Single-price mechanisms.

Consider the direct mechanism:

$$p^1(s,b) = \begin{cases} 1 \\ 0 \end{cases} \quad x^1(s,b) = \begin{cases} \pi & \text{if } 0 \leq s \leq \pi \text{ and } \pi \leq b \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

These mechanisms can be thought of as arising from a static game in which the seller and buyer announce simultaneous price bids. If the seller's bid is lower than the buyer's, then the buyer obtains the object at the average of the two bids. It is easy to see that one type of Nash equilibrium of this game (there exist others: see Chatterjee and Samuelson (1983), Leininger, Linhart and Radner (1987), and Satterthwaite and Williams (1988a,b)) has: the seller bid π , if $s \leq \pi$, and 1, otherwise; and the buyer bid π , if $b \geq \pi$, and 0, otherwise (for any π satisfying $0 < \pi < 1$). Observe that, given this strategy of the seller, knowledge of the seller's type does not help the buyer in formulating his own optimal bid.

Example 2: The monopoly mechanism.

Consider the mechanism:

$$p^2(s,b) = \begin{cases} 1 \\ 0 \end{cases} \quad x^2(s,b) = \begin{cases} g(s) & b \geq g(s), \\ 0 & \text{otherwise,} \end{cases}$$

where $g(s) \in \arg \max_{\pi} \{(\pi - s)(1 - F_2(\pi))\}$.

This type of mechanism was studied by Chatterjee and Samuelson (1983) for the case where $F_2(b) = b$. It corresponds to a game in which the seller has the ability to commit herself to a take-it-or-leave-it offer. Obviously, the optimal strategy for each seller is to name one of her monopoly prices.

Example 3: Revelation over time.

Assume that F_2 is absolutely continuous, and that $b - (1 - F_2(b))/f_2(b)$ is strictly increasing, so that the profit function to any seller $s \in [0,1]$, $W(\pi;s) = (\pi - s)(1 - F_2(\pi))$, is strictly quasiconcave in π . Let $g(s) = \arg \max_{\pi} W(\pi;s)$, and let $q(s)$ be any increasing C^1 function defined on $[0,1]$. Also assume that $g(s) < q(s) \leq 1$, that $q(0) < 1$, and that $q'(s) > 0$ whenever $q(s) < 1$. Consider the following mechanism:

$$p^3(s,b) = \begin{cases} h(s) \\ 0 \end{cases} \quad x^3(s,b) = \begin{cases} h(s)q(s) & b \geq q(s), \\ 0 & \text{otherwise.} \end{cases}$$

where $h(s) = \exp\{-\int_s^{\hat{s}} q'(s)[f_2(q(s))/(1 - F_2(q(s))) - 1/(q(s) - s)]ds\}$, and $\hat{s} = \sup\{s: q(s) < 1\}$. It is possible, though somewhat tedious, to check that $\{p^3, x^3\}$ is an ICBM. Observe that $h(\hat{s}) = 1$, and that $h(\cdot)$ is an

increasing function of s . Thus, if we convert $\{p^3, x^3\}$ to times and prices at which trade takes place (using transformation (4)), we see that high valuation sellers trade earlier, and at higher prices than low valuation sellers. As in Cramton (1984), time acts as a screening device. As discussed earlier, however, Cramton has low valuation sellers reveal before high valuation sellers reveal. We would obtain a similar revelation pattern here if we chose $0 \leq q(s) < g(s)$, with $q(1) > 0$. In that case the formula for the probability of trade would become:

$$h(s) = \exp\left\{\int_0^s q'(s) \left[\frac{f_2(q(s))}{1 - F_2(q(s))} - \frac{1}{q(s) - s} \right] ds\right\}.$$

Observe that in this mechanism, as in the previous one, the first serious price offered by the seller is perfectly revealing, and hence that the buyer receives no additional information when told what seller type he is facing. Observe also that we could have the seller offer declining price paths after revealing her type through her first serious offer; the explicit description of the resulting mechanism would become somewhat cumbersome.

Although a seller-first mechanism yields a price path for every seller, its requirements are insufficient to guarantee implementability. All that is assured is: if a seller is required to select from the menu of price paths implicit in the mechanism and is bound to follow her selection for all future time, then the seller truthfully reveals her type; furthermore, buyers accept or reject as posited by the mechanism. However, in the seller-offer bargaining game, a seller is free to initially follow a price path from the menu but to subsequently deviate in any number of arbitrarily

complicated ways, and thereby raise her profit above the value indicated by the static mechanism. Many such deviations may be deterred by adverse inferences, but it is never possible to deter a deviation which entails zero probability of future trade, i.e., a deviation which is tantamount to a unilateral refusal to further trade. Consider, for example, the following mechanism:

$$p^4(s,b) = \begin{cases} 1 & \\ 1/2 & \\ 16/27 & \\ 0 & \\ \end{cases} \quad x^4(s,b) = \begin{cases} 1/2 & 0 \leq s < 1/4 \text{ and } 3/4 < b \leq 1, \\ 1/8 & 0 \leq s < 1/4 \text{ and } 1/4 \leq b \leq 3/4, \\ 7/27 & \text{if } 1/4 \leq s \leq 7/16 \text{ and } 7/16 \leq b \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

One way to think of mechanism $\{p^4, x^4\}$ is that sellers in $[0, 1/4)$ pool by initially charging a price of $1/2$ and subsequently, at a time they discount to $1/2$, cut their price to $1/4$. Sellers in $[1/4, 7/16]$ also pool, and charge a price of $7/16$ at a time they discount to $16/27$, and never cut their price thereafter. Finally, sellers in $(7/16, 1]$ pool and charge a price of 1 forever. It is easy to verify that $\{p^4, x^4\}$ is an ICBM, and that it satisfies the additional requirements of a seller-first mechanism. However, the mechanism is not implementable by sequential equilibria in the seller-offer game. To see this, note that sellers in $[1/4, 7/16]$ could select to initially pool with sellers in $[0, 1/4)$, but subsequently refuse to cut their price any further. While it is entirely possible that this refusal would not lead to any further sales (since buyers expect prices to drop to $1/4$), sellers would be better off by following this strategy than by selecting the price path required by the static mechanism. The easiest way to see this is

to observe that the seller of type $7/16$ would make exactly zero profits from charging a price of $7/16$ at a time discounted to $16/27$, but would make strictly positive profits from charging a price of $1/2$ at time 0. Similarly, all sellers in $(1/4, 7/16]$ would prefer to deviate in the manner described above.

More generally, consider a seller of type s , who initially selected a price path geared toward type $s' \neq s$. A refusal to deal would then become profitable, relative to continued adherence to the price path of s' , when the price dropped below s . A stopping rule of this form would assure s a payoff of:

$$(6) \quad \tilde{U}_1(s', s) \equiv \int_{\hat{b}(s', s)}^1 [x(s', b) - sp(s', b)] d\mu_2(b),$$

where $\hat{b}(s', s) = \inf \{b \in [0, 1]: p(s', b) > 0, \text{ and } x(s', b)/p(s', b) \geq s\}$ denotes the last buyer to purchase from seller s' at a price at least equal to s . Observe that $\tilde{U}_1(s, s) = U_1(s, s)$ in seller-first mechanisms, but that $\tilde{U}_1(s', s) \geq U_1(s', s)$ generally. Thus, a seller-first mechanism must satisfy an additional incentive constraint to become implementable. We refer to seller-first mechanisms satisfying this additional incentive constraint as sequentially seller-first mechanisms.

Definition 2: A sequentially seller-first mechanism is a seller-first mechanism satisfying:

$$\tilde{U}_1(s, s) \geq \tilde{U}_1(s', s), \quad \forall s', s \in [0, 1].$$

Although a formal definition of implementation must wait until the next section, we conclude with a statement of Theorem 1, which follows from Lemma 2 and the above reasoning:

Theorem 1: Every seller-first mechanism which is not sequentially seller-first cannot be implemented by sequential equilibria in the seller-offer game.¹⁰

Finally, it is worth noting that there is an easily described subclass of the class of seller-first mechanisms for which $\tilde{U}_1(s',s) \equiv U_1(s',s)$: namely, mechanisms for which $p(s,b)$ takes on at most one value different from zero, for every $s \in [0,1]$. In such mechanisms, the seller never cuts her price below her first serious offer.

5. Implementability in the Seller-Offer Game

We will now make precise what we mean by the statement " $\{p,x\}$ is implementable by sequential equilibria of the seller-offer game." The most obvious definition to propose is that $\{p,x\}$ is implemented by a sequential equilibrium σ in the seller-offer game, with time interval z , when σ induces the mechanism $\{p,x\}$ through the transformation $p(s,b) = \exp\{-rt(s,b)\}$ and $x(s,b) = \pi(s,b)p(s,b)$, where t and π denote the times and prices at which s and b trade in the equilibrium. Such a definition is somewhat unsatisfactory for two reasons. First, since we are working with a discrete-time model, the set of implementable equilibria would most likely not be invariant with respect to the time interval between offers. A second (and somewhat related) reason is that we are interested in understanding the

outcome of the bargaining situation when there is virtually no restraint to the rate at which the seller can make offers. These considerations lead to the following definition:

Definition 3: Let $\{p,x\}$ be an arbitrary ICBM. We will say that $\{p,x\}$ is implementable by sequential equilibria of the seller-offer game if there exists a sequence $\{\sigma^n, z^n\}_{n=1}^{\infty}$ such that:

- (i) $z^n \downarrow 0$ and, for every $n \geq 1$, σ^n is a sequential equilibrium of the seller-offer game where the time between offers is z^n ; and
- (ii) for each $\epsilon > 0$, and for each $\mu \in \{\mu_1 \times \mathcal{Q}, \mathcal{Q} \times \mu_2\}$, where \mathcal{Q} denotes the Lebesgue measure on $[0,1]$, the ICBM's $\{p^n, x^n\}$ induced by σ^n satisfy:

$$(7) \quad \mu\{(s,b): |p^n(s,b) - p(s,b)| > \epsilon\} \rightarrow 0$$

$$(8) \quad \mu\{(s,b): |x^n(s,b) - x(s,b)| > \epsilon\} \rightarrow 0.$$

While our notion of convergence of $\{p^n, x^n\}$ to $\{p,x\}$ is fairly weak, namely convergence in $(\mu_1 \times \mathcal{Q})$ - and $(\mathcal{Q} \times \mu_2)$ -measure, its implications are very strong. For, as we show in Proposition 1 below, (7) and (8) imply uniform convergence of buyer and seller utilities.

Proposition 1: Let $\{p,x\}$ be an arbitrary ICBM, and let $\{p^n, x^n\}$ be a sequence of ICBM's such that (7) and (8) hold. Then, if we let $U_1^n(s,s) = \bar{x}_1^n(s) - s\bar{p}_1^n(s)$ and $U_2^n(b,b) = b\bar{p}_2^n(b) - \bar{x}_2^n(b)$, we have:

$$|U_1^n(s,s) - U_1(s,s)| \rightarrow 0, \text{ uniformly in } s; \text{ and}$$

$$|U_2^n(b,b) - U_2(b,b)| \rightarrow 0, \text{ uniformly in } b.$$

Proof: Since $\{p,x\}$ is an ICBM, Myerson and Satterthwaite (1983, equation 4) implies that $U_1(s,s) = \int_s^1 \bar{p}_1(v_1) dv_1$. An analogous formula holds for $U_2^n(s,s)$. Consequently, for all $s \in [0,1]$:

$$|U_1^n(s,s) - U_1(s,s)| \leq \int_s^1 \int_0^1 |p^n(v_1, v_2) - p(v_1, v_2)| d\mu_2(v_2) dv_1.$$

Since $p^n \rightarrow p$ in $(\mathcal{Q} \times \mu_2)$ -measure, this implies that $\forall \epsilon > 0, \exists \bar{n}(\epsilon)$ such that $\forall n \geq \bar{n}: |U_1^n(s,s) - U_1(s,s)| < \epsilon$, for all $s \in [0,1]$. Also, since $U_2(b,b) = \int_0^b \bar{p}_2(v_2) dv_2$, and since $p^n \rightarrow p$ in $(\mu_1 \times \mathcal{Q})$ -measure, the proof that $|U_2^n(b,b) - U_2(b,b)| \rightarrow 0$, uniformly in b , proceeds entirely analogously. \square

One remark is in order here. If we are merely interested in convergence of traders' utilities, we do not need to require (8) in the definition of implementability. This motivates the next definition:

Definition 4: $\{p\}$ is implementable if there exists an x such that $\{p,x\}$ is implementable.

6. Stationary Sequential Equilibria and the No-Trade Theorem

Before delving into the statement of our main theorem, it is necessary to be a little more precise about the nature of the strategy spaces and the meaning of sequential equilibrium. At the start of each period t , the history of the game is a sequence of rejected price offers: $H^t = \{p^0, p^1, \dots, p^{t-1}\}$. Given this history, the seller updates her belief about

the buyer type to $\mu_s^t(b|H^t)$. The seller then makes a price offer p^t according to her own type and the history, using the strategy:

$$\hat{s}^t: [0,1] \times H^t \rightarrow [0,1].$$

Having observed history H^t and the seller's offer of p^t in period t , the buyer forms a posterior conjecture, $\mu_b^t(s|H^t, p^t)$, of the seller's valuation. He then decides whether or not to accept the offer according to his own type, the pre-period history and the current price, using the strategy:

$$\hat{b}^t: [0,1] \times H^t \times [0,1] \rightarrow \{\text{accept, reject}\}.$$

A sequential equilibrium is a pair of strategy profiles $\{\hat{s}^t, \hat{b}^t\}_{t=0}^{\infty}$ and a system of beliefs $\{\mu_s^t(b|H^t), \mu_b^t(s|H^t, p^t)\}_{t=0}^{\infty}$ that satisfy the requirements of sequential rationality and consistency. Consistency says that the posterior beliefs conditioned on the history H^t must be computed through Bayes' rule whenever possible. Sequential rationality says that given any history and the induced posterior belief, strategies from then on must be optimal for both players.

In analyzing sequential equilibria of bargaining games with incomplete information, several authors (Gul, Sonnenschein and Wilson, 1986; Gul and Sonnenschein, 1988; Cho, 1988) have found it useful to impose a stationarity restriction on strategies and updating rules. Stationarity is attractive on simplicity grounds, since in stationary equilibria players do not condition their actions on the entire history of the game, but only on a lower-dimensional summary of it (state space). Formally, we may define

stationarity in the seller-offer game as follows:¹¹

Definition 5: A strong stationary sequential equilibrium in the seller-offer game is a sequential equilibrium satisfying the following two requirements:

- (i) players' (current-period) strategies are constant on histories H_1 and H_2 that induce the same beliefs, i.e.,

$$\begin{aligned} \mu_b^{t-1}(s|H_1^{t-1}, p_1^{t-1}) &= \mu_b^{\tau-1}(s|H_2^{\tau-1}, p_2^{\tau-1}) \text{ and } \mu_s^t(b|H_1^t) = \mu_s^\tau(b|H_2^\tau) \\ \Rightarrow \hat{s}^t(s, H_1^t) &= \hat{s}^\tau(s, H_2^\tau), \text{ and} \end{aligned}$$

$$\begin{aligned} \mu_s^t(b|H_1^t) &= \mu_s^\tau(b|H_2^\tau) \text{ and } \mu_b^t(s|H_1^t, p) = \mu_b^\tau(s|H_2^\tau, p) \\ \Rightarrow \hat{b}^t(b, H_1^t, p) &= \hat{b}^\tau(b, H_2^\tau, p); \text{ and} \end{aligned}$$

- (ii) players' updating rules are constant on histories H_1 and H_2 that induce the same beliefs, i.e.,¹²

$$\begin{aligned} \mu_b^{t-1}(s|H_1^{t-1}, p_1^{t-1}) &= \mu_b^{\tau-1}(s|H_2^{\tau-1}, p_2^{\tau-1}) \text{ and } \mu_s^t(b|H_1^t) = \mu_s^\tau(b|H_2^\tau) \\ \Rightarrow \mu_b^t(s|H_1^t, p) &= \mu_b^\tau(s|H_2^\tau, p), \forall p, \text{ and} \end{aligned}$$

$$\begin{aligned} \mu_s^t(b|H_1^t) &= \mu_s^\tau(b|H_2^\tau) \text{ and } \mu_b^t(s|H_1^t, p) = \mu_b^\tau(s|H_2^\tau, p) \\ \Rightarrow \mu_s^{t+1}(b|H_1^{t+1}) &= \mu_s^{\tau+1}(b|H_2^{\tau+1}). \end{aligned}$$

As in the case of one-sided incomplete information, it is often necessary to slightly relax the requirement on the seller's strategy in order to ensure

the existence of stationary equilibria. In a weak stationary sequential equilibrium the seller's strategy is allowed to depend not only on current beliefs, but also on the previous price she charged:

$$\begin{aligned} \mu_b^{t-1}(s|H_1^{t-1}, p_1^{t-1}) &= \mu_b^{\tau-1}(s|H_2^{\tau-1}, p_2^{\tau-1}), \quad \mu_s^t(b|H_1^t) = \mu_s^\tau(b|H_2^\tau), \\ \text{and } p_1^{t-1} &= p_2^{\tau-1} \Rightarrow \hat{s}^t(s, H_1^t) = \hat{s}^\tau(s, H_2^\tau). \end{aligned}$$

Of the two papers on the seller-offer game discussed in Section 2, only Cho's (1988) equilibrium satisfies the stationarity restrictions. The revelation-over-time equilibrium exhibited by Cramton (1984) is not stationary: it satisfies neither property (i) nor property (ii) in Definition 5. It is perhaps worth elaborating on exactly where the nonstationarity is located in the Cramton equilibrium. Let us consider two different histories: in H_1^n , the seller has mimicked type s for the first n periods of the game; in H_2^{n+1} , the seller has mimicked type s for n periods and charged a price of 1 in period n . Clearly, no buyer will buy in period n of H_2^{n+1} as the price then charged exceeds his valuation. Consequently, the seller will have the same beliefs about the buyer entering period $(n + 1)$ after H_2^{n+1} as she does entering period n after H_1^n . In addition, Cramton requires that the buyer's conjecture conditional on a price higher than the equilibrium price remains the same as the conjecture on the equilibrium path. Thus, entering both period n after H_1^n and period $(n + 1)$ after H_2^{n+1} , the beliefs held by the buyer and the seller are identical. However, in the first case the seller strategy specifies charging a price of $p_n(s)$, while in the second case it calls for a price of $p_{n+1}(s) < p_n(s)$ (where $p_j(s)$ denotes the price s charges along the equilibrium path in

period $j = n, n + 1$), thus violating (i). In addition, as emphasized by Wilson (1987, p. 48), the buyer's updating rule is nonstationary: if after H_1^n the seller made an offer of $p_{n+1}(s) < p_n(s)$, the buyer would revise his beliefs downward, whereas no such updating occurs when the same price is charged in period $(n + 1)$ after H_2^{n+1} .

The two equilibria nevertheless do have important properties in common. First, in both equilibria, the seller reveals her type before any trade takes place. In Cho (1988), the separation occurs entirely in the first period;¹³ in Cramton (1984), the separation occurs over time. Second, both equilibria specify a strategy for the zero seller type (and associated buyer beliefs) that implies the celebrated Coase conjecture: as the time between successive offers shrinks to zero, the zero seller's initial price converges to zero. It remains an open question today whether stationarity (as stated in Definition 5) together with revelation-before-trade (at least for the zero seller type) are sufficient to imply the Coase conjecture property for the zero seller's strategy.¹⁴ If, on the other hand, once the zero seller has revealed, buyers never update beliefs, as in Cramton (1984), and if the zero seller's price is monotone in the state (the highest remaining buyer valuation), then it is easy to show that the Coase conjecture must hold. However, as we show below, this leads to disastrous results: as the time between successive offers shrinks to zero, the induced mechanism approaches the zero-trade mechanism (i.e., $p(\cdot, \cdot) \equiv 0$; $x(\cdot, \cdot) \equiv 0$).¹⁵

Lemma 3: Let the distribution function $F_1(\cdot)$ satisfy $F_1(0) = 0$. Then there exists an absolutely continuous distribution function $G(\cdot)$ such that $F_1(s) \leq G(s)$, for all $s \in [0, 1]$, and $G(0) = 0$.

Proof: Let $A = \{(x,y): x \in [0,1] \text{ and } 0 \leq y \leq F_1(x)\}$, and let B be the convex hull of A . Also define $G(x) = \sup\{y: (x,y) \in B\}$. It is straightforward to verify that $G(\cdot)$ is a continuous increasing concave function on $[0,1]$, and that $G(0) = 0$ and $G(1) = 1$. The concavity of $G(\cdot)$ then establishes its absolute continuity. []

Theorem 2: Consider any sequence $\{\sigma^n, z^n\}_{n=1}^{\infty}$ of sequential equilibria of the seller-offer game and associated time intervals between successive offers such that:

- (i) $z^n \downarrow 0$; and
- (ii) the Coase conjecture holds for the zero seller type.

Also, suppose that the distribution of seller types satisfies $F_1(0) = 0$. Then the aggregate probability of trade, $\int_0^1 \int_0^1 p^n(s,b) d\mu_2(b) d\mu_1(s)$, converges to zero as $n \rightarrow \infty$.

Proof: Define $G(\cdot)$ as in Lemma 3. Since $\bar{p}_1(\cdot)$ is monotone:

$$(9) \quad \int_0^1 \bar{p}_1^{-n}(s) d\mu_1(s) \leq \int_0^1 \bar{p}_1^{-n}(s) g(s) ds.$$

where $g(\cdot)$ is the density of $G(\cdot)$. By the Coase conjecture property, $U_1^n(0,0) = \int_0^1 \bar{p}_1^{-n}(s) ds \rightarrow 0$ as $n \rightarrow \infty$. The absolute continuity of $G(\cdot)$ then implies that the right side of (9) goes to zero as well, and hence the desired result follows. []

It is easy to understand what economic forces drive Theorem 2: if the zero seller's profits were approximately equal to zero, but a higher seller also had substantial probability of trade, the latter would be imitated. Incentive compatibility thus forces a low aggregate probability of trade.

In order to avoid the disastrous consequences of Theorem 2, it is necessary to discard at least one of the four joint properties implying the Coase conjecture (stationarity, revelation-before-trade, monotonicity and absorbing beliefs at $s = 0$). We have chosen to drop full revelation at the lower end of the support of seller types: all our equilibria, while stationary, involve some amount of pooling of low seller types with the seller of type zero.¹⁶

7. The Main Result

In this section we prove Theorem 3, which establishes the implementability of sequentially seller-first mechanisms. Its proof will require two sets of conditions: some mild restrictions on the distribution functions of player valuations, and some technical assumptions on the mechanisms which will be implemented.

Assumption 1: $F_1(\cdot)$ and $F_2(\cdot)$ satisfy the following properties:

- (i) the support of each distribution function equals $[0,1]$;
- (ii) $F_1(\cdot)$ is absolutely continuous on $[0,\epsilon]$, for some $\epsilon > 0$; and
- (iii) $F_2(\cdot)$ is absolutely continuous on $[0,1]$, and there exist L, M ($0 < M \leq L < \infty$) and $\alpha > 0$ such that $Mx^\alpha \leq F_2(x) \leq Lx^\alpha$ for all

$x \in [0,1]$.

Assumption 2: $\{p,x\}$ satisfies the following conditions:

- (i) $\tilde{U}_1(s',s)$ is quasiconcave in s' for all s ;¹⁷
- (ii) there exists $K < \infty$ such that $p(\cdot,b)$ has total variation less than K for almost every $b \in [0,1]$; and
- (iii) for every $\epsilon > 0$, there exists $\delta(\epsilon) > 0$ such that $\lim_{\epsilon \downarrow 0} \delta(\epsilon) = 0$ and $p(s,s + \epsilon) = 0$ for all $s \in [0, \hat{s} - \delta(\epsilon)]$, where $\hat{s} = \sup\{s: \bar{p}_1(s) > 0\}$.

We now state the main theorem:

Theorem 3: Suppose that the distribution functions satisfy Assumption 1 and let $\{p,x\}$ be any sequentially seller-first bargaining mechanism which satisfies Assumption 2. Then $\{p,x\}$ is implementable by stationary sequential equilibria in the seller-offer game.

The proof of Theorem 3 is long and intricate, and is therefore relegated to the Appendix. In the remainder of this section, we will highlight some of the ideas and constructions in the general proof by demonstrating the implementability of our first two examples of seller-first mechanisms.

First we will directly argue the implementability of single-price mechanisms (Example 1 of Section 4). For any single price π ($0 < \pi < 1$), the transformation (4) suggests a price path $\pi_0(t) \equiv \pi$ for all seller types

$s \in [0, \pi]$. Meanwhile, seller types $s \in (\pi, 1]$ make no sales, suggesting a price path $\pi_1(t) \equiv 1$. This naive attempt at implementation runs up against immediate difficulties, as the zero seller type faces inexorable temptations to cut her price below π after she has sold to all buyer valuations above π . Indeed, in any sequential equilibrium, the zero seller type must eventually transact with every potential buyer (possessing positive valuation). Sequentiality thus mandates that the zero type's price path converges toward zero. Simplicity recommends that we utilize a single, smoothly-descending price path for all types $s \in [0, \pi]$, who remain pooled so long as it stays profitable.

We redefine the lower price path by $\pi_0(t) = e^{-\lambda t} \pi$, and define $\pi_1(t)$ as before. In equilibrium, a seller with valuation s charges $\pi_0(t)$ in all periods such that $\pi_0(t)$ exceeds her valuation s and charges $\pi_1(t)$ otherwise. If the buyer observes a history inconsistent with this strategy, his expectations change in two respects: (a) he comes to believe that the seller's valuation equals zero, and never updates his beliefs thereafter; and (b) he expects that the seller will follow a "Coase conjecture" price path in which offers are exceedingly low (compared to the highest buyer valuation remaining in the market).

Why does the seller adhere to her equilibrium strategy at all moments in time? (This is demonstrated formally in part III of the proof.) The answer has two parts, depending on whether the seller's valuation is relatively closer to zero or to the current price $\pi_0(t)$. If the seller's valuation is in the vicinity of zero, then the seller retains a tangible probability of a relatively profitable sale by continuing to adhere to $\pi_0(\cdot)$, whereas a deviation would trigger the Coase conjecture. Meanwhile,

if the seller's valuation is in the vicinity of the current price, then while continuation profits are insubstantial, a deviation would trigger adverse inferences, which choke off profitable sales altogether.

To demonstrate the implementability of the single-price mechanism, we now consider a sequence of equilibria of the above form, using rates of descent $\lambda^n \downarrow 0$. Observe, for arbitrarily small λ^n , that arbitrarily close to all buyer valuations in $[\pi, 1]$ accept the initial offer $\pi_0(0)$, and that arbitrarily close to all buyer valuations in $[0, \pi)$ defer purchase until arbitrarily far into the future. Finally, as $\lambda^n \downarrow 0$, seller types "drop off" the price path $\pi_0(\cdot)$ arbitrarily slowly. Thus, the probabilities of trade implicit in the sequence of equilibria converge in measure to $\{p^1(\cdot, \cdot), x^1(\cdot, \cdot)\}$, establishing implementability.

It should further be observed that these equilibria are stationary, as: (a) the state evolves in every period along the equilibrium path; and (b) behavior after a deviation is also stationary (see Fudenberg, Levine and Tirole, 1985; Gul, Sonnenschein and Wilson, 1986). Moreover, these equilibria satisfy even a stronger sense of stationarity than we require in Definition 5; the buyer's acceptance behavior depends only on his beliefs about the seller and the current seller offer.¹⁸

Let us also note that this is not literally the way in which our general proof of Theorem 3 treats a single-price mechanism. (The general proof defines multiple pools of seller types, while that is clearly unnecessary to handle a single-price mechanism.) However, we literally do treat the "bottom pool" (which includes the zero type) in exactly the above manner.

Now we will directly argue the implementability of the monopoly

mechanism (Example 2 of Section 4). Define a grid of equidistant seller types such that $0 < s_1 < s_2 < \dots < s_N < 1$. Begin a menu of price paths by defining $\pi_k(t) \equiv g(s_k)$, for $k = 1, \dots, N$, where $g(s_k)$ denotes a static monopoly price of s_k . Enlarge the menu of price paths by also offering the set of constant price paths $\pi(t) \equiv q$, for all $q \in [g(s_N), 1]$. Consider an auxiliary game in which the seller must select one price path from this menu (and becomes committed to following it forever), and the buyer then optimizes against her selection. Observe that this game has a unique sequential equilibrium, which induces a seller-first mechanism. It has the property that seller types partition themselves into $(N + 1)$ intervals such that the k th interval (which contains the type s_k) charges the monopoly price of s_k , for $k = 1, \dots, N$, and the last interval $(s_N, 1]$ fully separates by each type charging her own monopoly price.

As argued above, to make this a sequential equilibrium, we should add one additional price path, $\pi_0(t) \equiv e^{-\lambda t} g(0)$ and invoke a change in buyer expectations after a detectable seller deviation from any of the aforementioned price paths. To demonstrate implementation, we merely need to let $\lambda^n \downarrow 0$ and $N^n \uparrow \infty$.

The implementation of the single-price and monopoly mechanisms was straightforward on account that their implicit price paths were constant over time. In general, however, the price paths $\{\pi_k(t)\}_{k=1}^N$ from seller-first mechanisms are nonconstant, leading to some complications. First, for each seller type on the grid, we will need to take a discrete approximation to her continuous-time price path. Second, $\tilde{U}_1(s', s)$ will generally not coincide with $U_1(s', s)$, so it will require a delicate argument to show that the separation of seller types in the auxiliary game is preserved when the

seller has the option to depart from the chosen price path.

Third, the profit calculations in the proof require that (except for the bottom interval of seller types) each seller type selects a price path which, once chosen, is strictly adhered to (i.e., the seller is not allowed to drop off the price path). Obviously, when the price paths are constant over time, there is no incentive problem in assuring permanent adherence. In general, however, we must be careful to argue that the most preferred price path of each seller type (outside the bottom interval) is one which never dips below her valuation. This is particularly a problem for seller types who have zero probability of making sales in the mechanism being implemented. Our purpose in introducing a continuum of constant price paths taking values near one is to provide most-preferred paths for high-valuation sellers to which they permanently adhere.

Fourth, the general specification of price path $\pi_0(\cdot)$ requires the property that it always asymptotically approach zero, in order to permit the zero seller to eventually transact with all positive-valuation buyer types. At the same time, it should be defined in such a way that seller types in the interval $[0, s_1)$ nontrivially partition themselves between selecting paths $\pi_0(\cdot)$ and $\pi_1(\cdot)$. Informally speaking, we construct $\pi_0(\cdot)$ so as to begin by approximating the zero type's price path from the mechanism yet to terminate with an exponential rate of descent toward zero. However, the formal proof requires a fixed-point argument: seller types allocate themselves between $\pi_0(\cdot)$ and $\pi_1(\cdot)$ according to the profits they expect; but the profits attributable to selecting $\pi_0(\cdot)$ depend on which seller types choose that price path (since this, in turn, influences optimal buyer purchase behavior). Additional complications of a less intuitive nature are also

handled in the proof.

8. The Efficiency of Monopoly

An interesting question one might ask is: What is the most efficient sequentially seller-first mechanism, and how does its efficiency compare to that of the most efficient ICBM? In this section we will show that if the efficiency objective consists of maximizing expected gains from trade, and provided the distribution functions satisfy some regularity conditions, the answer to the first question is surprisingly simple. We will also argue that the most efficient sequentially seller-first mechanism may realize a level of gains from trade that comes close to that of the ex ante efficient ICBM.

Maximizing efficiency over the set of sequentially seller-first mechanisms is a potentially daunting task. Lemma 4 below greatly simplifies the search for such a mechanism. However, before stating the lemma, we must introduce a definition.

Definition 6: A mechanism $\{p, x\}$ is a 0-1 mechanism if there exists a nondecreasing function $\phi(s)$ such that $x(s, b) = p(s, b) = 0$ for $0 \leq b < \phi(s)$ and $p(s, b) = 1$ for $\phi(s) \leq b \leq 1$.

Lemma 4: Suppose $0 \leq p(s, b) \leq 1$ maximizes:

$$(10) \quad f(p) \equiv \int_0^1 \int_0^1 (v_2 - v_1) p(v_1, v_2) d\mu_1(v_1) d\mu_2(v_2),$$

subject to the constraint:

$$(11) \quad h(p,s) \equiv \int_0^1 t_2 p(s,t_2) d\mu_2(t_2) - \int_0^1 \int_0^{v_2} p(s,t_2) dt_2 d\mu_2(v_2) \\ - \int_s^1 \bar{p}_1(t_1) dt_1 - s\bar{p}_1(s) = 0, \text{ for all } s \in [0,1].$$

Suppose also that $p(\cdot, \cdot)$ satisfies the requirements of a 0-1 mechanism. Then, letting $x(s,b) = \int_0^b r dp_s(r)$, the mechanism $\{p,x\}$ maximizes the gains from trade among all sequentially seller-first mechanisms.

Proof: We will show that $p(\cdot, \cdot)$ maximizes the gains from trade over all seller-first mechanisms. Since $\{p,x\}$ is a 0-1 mechanism, it is also sequentially seller-first and hence maximizes the expected gains from trade over this smaller set as well.

To prove the first claim, observe that in any ICBM:

$$\bar{x}_1(s) = U_1(s,s) + s\bar{p}_1(s) = \int_s^1 \bar{p}_1(t_1) dt_1 + s\bar{p}_1(s).$$

Furthermore, in a seller-first mechanism, $x(s,b) = \int_0^b r dp_s(r) = bp(s,b) - \int_0^b p(s,r) dr$. Hence we have $\bar{x}_1(s) = \int_0^1 t_2 p(s,t_2) d\mu_2(t_2) - \int_0^1 \int_0^{v_2} p(s,t_2) dt_2 d\mu_2(v_2)$. Thus, in any seller-first mechanism, constraint (11) must hold.

Observe now that $\{p,x\}$ is an ICBM (it is a 0-1 mechanism), and that it satisfies the requirements of a seller-first mechanism. []

Before stating our efficiency theorem, let us first make our "regularity" assumptions on the distribution functions F_1 and F_2 :¹⁹

Assumption 3: $F_1(\cdot)$ is convex and C^2 . $F_2(\cdot)$ is concave and C^2 .

Furthermore, $v_2 - (1 - F_2(v_2))/f_2(v_2)$ is an increasing function.

Theorem 4: Suppose F_1 and F_2 satisfy Assumption 3. Then the most efficient sequentially seller-first mechanism is the monopoly mechanism.

Proof: We will show that there exists a function $\rho(s)$ of bounded variation on $[0,1]$ such that the 0-1 mechanism with boundary $g(s)$ satisfying:

$$g(s) - s = (1 - F_2(g(s)))/f_2(g(s)) = 0$$

maximizes the Lagrangian $L(p, \rho) = f(p) + \int_0^1 h(p, t_1) d\rho(t_1)$ subject to $0 \leq p(\cdot, \cdot) \leq 1$. The sufficiency theorem of Luenberger (1969, p. 220) then implies that $p(\cdot, \cdot)$ solves the infinite-dimensional linear program (10)-(11). Observe first that we may rewrite L as:

$$L = \int_0^1 \left\{ \int_0^1 [v_2 - v_1 - \rho(v_1)/f_1(v_1)] p(v_1, v_2) d\mu_1(v_1) \right. \\ \left. + \int_0^1 [v_2 - v_1 - (1 - F_2(v_2))/f_2(v_2)] p(v_1, v_2) d\rho(v_1) \right\} d\mu_2(v_2).$$

Now let $\rho(v_1) = f_1(v_1)[g(v_1) - v_1]$. Then the coefficient of $p(v_1, v_2)$ under L becomes:

$$(12) \quad \{v_2 - g(v_1) + [v_2 - v_1 - (1 - F_2(v_2))/f_2(v_2)] \\ \{ (f_1'(v_1)/f_1(v_1))(g(v_1) - v_1) + g'(v_1) - 1 \} \}.$$

Notice that, evaluated at the point $(v_1, g(v_1))$, the expression in (12) is

equal to zero. Next, we will argue that its derivative with respect to v_2 is nonnegative. This will imply that the mechanism described above maximizes L , and hence establish its optimality. Now, the derivative of (12) is given by

$$(13) \quad 1 + (2 + \omega(v_2))[(f'_1(v_1)/f_1(v_1))(g(v_1) - v_1) + g'(v_1) - 1],$$

where $\omega(v_2)$ is implicitly derived from $(d/dv_2)[v_2 - (1 - F_2(v_2))/f_2(v_2)] \equiv 1 + \omega(v_2) > 0$. Concavity of F_2 implies that $\omega(v_2) \leq 0$. Also, convexity of F_1 and $g(v_1) \geq v_1$ implies that (13) $\geq 1 + (2 + \omega(v_2))(g'(v_1) - 1)$. Since $g'(v_1) = 1/(2 + \omega(g(v_1)))$, and since $-2 < \omega(v_2) \leq 0$, we may bound the latter expression again to obtain (13) $> 1 + (2 + (\omega(v_2))(-1/2)) = -\omega(v_2)/2 \geq 0$. This establishes the desired result. []

One immediate consequence of Theorem 4 is that when F_1 and F_2 satisfy Assumption 3, then the optimal seller-first mechanism is implementable. In fact, this mechanism fares quite well in terms of efficiency when compared to the ex-ante efficient ICBM. When F_1 and F_2 are linear, for example, Myerson and Satterthwaite (1983) established the optimality of the mechanism:

$$p^5(s, b) = \begin{cases} 1 \\ 0 \end{cases} \quad x^5(s, b) = \begin{cases} (s + b + 1/2)/3 & \text{if } b \geq s + 1/4 \\ 0 & \text{otherwise.} \end{cases}$$

Some simple computations show that the sum of the players' ex-ante expected

utility in that case is equal to $9/64$. The monopoly mechanism, on the other hand, yields a total utility of $1/8$. Thus only one-ninth of the feasible gains from trade are lost by restricting attention to seller-first mechanisms.

9. Conclusion

Consider the problem of a durable goods monopolist whose marginal cost is incompletely known by the outside world. Let us reinterpret our "single buyer with (potential) valuations" as a "continuum of buyers with (actual) valuations" distributed according to the function $F_2(\cdot)$. With this reinterpretation, we now face the question--analogous to the bilateral bargaining problem--of what will be the pricing behavior of a rational durable goods monopolist who faces rational (but incompletely informed) consumers. The foregoing game immediately becomes an attractive model of durable goods monopoly with private marginal cost.²⁰

Contra Coase's conjecture, we have established (in Theorem 3) that this durable goods monopoly game possesses stationary sequential equilibria in which the seller, ex ante, expects substantial profits. The intuition behind this result is quite natural: the consumers' incomplete information provides the monopolist with a credible means to hide her valuation and hence not to reduce her price. Were the monopolist to unexpectedly cut her price, consumers would view this as a sign that her marginal cost was low and her resolve was weak, and they would accordingly expect even further price slashing.

Two equilibria, in particular, are very easy to describe. The first comes arbitrarily close to the story that the monopolist truthfully reveals

her cost in the initial period and charges the static monopoly price (for her cost) forever, i.e., the monopoly mechanism. The second comes arbitrarily close to a description that the monopolist truthfully reveals in the initial period whether or not her cost exceeds some price π , and in the latter event she charges that price forever, i.e., a single-price mechanism. Both equilibria have the appealing feature of relative simplicity.

Furthermore, we have argued that even if proponents of the Coase conjecture were to win the battle, they would still lose the war.²¹ With private seller valuations, the Coase conjecture does not lead to an especially happy state of affairs. Suppose that there is no gap between the consumers' minimum-possible valuation and the monopolist's minimum-possible marginal cost. Then we have demonstrated (in Theorem 2) that, as the time interval between the monopolist's offers approaches zero, the expected quantity of sales (discounted for time of trade) must also converge to zero. The Coase conjecture chokes off the market: in order for the low-cost seller to earn minimal unit profits, the high-cost seller must make minimal overall sales.

Returning to the bargaining problem, we should remark that since our model involves two-sided incomplete information, all the results have obvious analogues in the game where the buyer makes all the offers. Clearly, single-price mechanisms are implementable in the buyer-offer game; also, the "monopsony mechanism" where the buyer fully reveals and charges his monopsony price is implementable. More generally, all sequentially buyer-first mechanisms are implementable in stationary sequential equilibria of the buyer-offer game, and the monopsony mechanism is optimal within the class of buyer-first mechanisms when $F_1(\cdot)$ is concave and $F_2(\cdot)$ is convex.

Finally, we hope that our article contributes some insight into the general nature of sequential equilibrium in games of incomplete information. It is easy to prematurely conclude that a folk theorem will hold in most sequential games, on the grounds that "any action can be deterred by beliefs." We have shown (in Theorem 1) that such reasoning is false in the seller-offer bargaining game with two-sided incomplete information.

It is impossible (within the rules of the game) to induce a seller to continue along a price path for which all future offers lie below her valuation and for which the probability of acceptance is positive. Independent of the buyer's inferences, a refusal to deal always dominates. More generally, folk theorems may fail in games of incomplete information on account that certain actions cannot be deterred by any conjectures, no matter how extreme or unreasonable. We will continue on this theme in future papers.

References

- Admati, A. and M. Perry (1987), "Strategic Delay in Bargaining," Review of Economic Studies, 54, 345-364.
- Ausubel, L. and R. Deneckere (1986), "Reputation in Bargaining and Durable Goods Monopoly," Technical Report No. 491, Institute for Mathematical Studies in the Social Sciences, Stanford University, revised 1988.
- Ausubel, L. and R. Deneckere (1988), "A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information," Journal of Economic Theory (forthcoming).
- Bulow, J. (1982), "Durable-Goods Monopolists," Journal of Political Economy, 90, 314-332.
- Chatterjee, K. and L. Samuelson (1986), "Bargaining Under Incomplete Information: The Unrestricted Offers Case," mimeo, Pennsylvania State University.
- Chatterjee, K. and L. Samuelson (1987), "Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers," The Review of Economic Studies, 54 (2), 175-192.
- Chatterjee, K. and W. Samuelson (1983), "Bargaining Under Incomplete Information," Operations Research, 31, 835-851.
- Cho, I. (1988), "Uncertainty and Delay in Bargaining: Old Question and New Approach," mimeo, The University of Chicago.
- Cramton, P. (1984), "Bargaining with Incomplete Information: An Infinite-Horizon Model with Continuous Uncertainty," Review of Economic Studies, 51, 579-593.
- Cramton, P. (1985), "Sequential Bargaining Mechanisms," in Game Theoretic

- Models of Bargaining, A. Roth (ed.), Cambridge University Press, London.
- Cramton, P. (1988), "Strategic Delay in Bargaining with Two-Sided Uncertainty," Journal of Economic Theory, forthcoming.
- Fudenberg, D., D. Levine and J. Tirole (1985), "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information," in Game Theoretic Models of Bargaining, op. cit., 73-98.
- Fudenberg, D. and J. Tirole (1983), "Sequential Bargaining with Incomplete Information," Review of Economic Studies, 50 (2), 221-247.
- Grossman, S. and M. Perry (1986), "Sequential Bargaining under Asymmetric Information," Journal of Economic Theory, 39, 119-154.
- Gul, F. and H. Sonnenschein (1988), "On Delay in Bargaining with One-Sided Uncertainty," Econometrica, 56 (3), 601-612.
- Gul, F., H. Sonnenschein and R. Wilson (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture," Journal of Economic Theory, 39, 155-190.
- Leininger, W., P. Linhart and R. Radner (1986), "The Sealed-Bid Mechanism for Bargaining with Incomplete Information," Discussion Paper No. A-95, Sonderforschungsbereich 303, University of Bonn.
- Myerson, R. (1985), "Analysis of Two Bargaining Problems with Incomplete Information," in Game Theoretic Models of Bargaining, op. cit., 115-147.
- Myerson, R. and M. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading," Journal of Economic Theory, 28, 265-281.
- Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," Econometrica, 50, 97-109.

- Rubinstein, A. (1985), "Choice of Conjectures in a Bargaining Game with Incomplete Information," in Game Theoretic Models of Bargaining, op. cit., 99-114.
- Rubinstein, A. (1987), "A Sequential Theory of Bargaining," in Advances in Economic Theory, T. Bewley (ed.), Cambridge University Press, London.
- Satterthwaite, M. and S. Williams (1988a), "Bilateral Trade with the Sealed Bid Double Auction: Existence and Efficiency," Journal of Economic Theory, forthcoming.
- Satterthwaite, M. and S. Williams (1988b), "The Role of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large," CMSEMS Discussion Paper No. 741, Northwestern University.
- Sobel, J. and I. Takahashi (1983), "A Multi-Stage Model of Bargaining," Review of Economic Studies, 50, 411-426.
- Stokey, N. (1981), "Rational Expectations and Durable Goods Pricing," Bell Journal of Economics, 12, 112-128.
- Williams, S. (1987), "Efficient Performance in Two Agent Bargaining," Journal of Economic Theory, 41, 154-172.
- Wilson, R. (1987), "Game-Theoretic Analyses of Trading Processes," in Advances in Economic Theory, op. cit., 33-70.

Notes

1. The mechanism design result is due to Myerson and Satterthwaite (1983).
2. In the seller-offer bargaining game with bilateral private information, we define the "state" to be: the seller's most recent beliefs about the buyer; the buyer's most recent beliefs about the seller; and the current (if any) offer on the table. See Section 6, below.
3. Authors who have studied stationarity in bargaining games include: Fudenberg, Levine and Tirole (1985); Gul, Sonnenschein and Wilson (1986); Wilson (1987); Gul and Sonnenschein (1988); Cho (1988); and Cramton (1988).
4. Obviously, the monopoly mechanism will be ex ante efficient under some choices of welfare weights--for example, when the objective is to maximize the expected seller surplus. The implementability of ex ante efficient mechanisms, more generally, will be addressed in a sequel to this paper.
5. Papers on the seller-offer bargaining game with one-sided incomplete information and the related problem of durable goods monopoly include: Bulow (1982); Stokey (1981); Fudenberg and Tirole (1983); Sobel and Takahashi (1983); Fudenberg, Levine and Tirole (1985); Gul, Sonnenschein and Wilson (1986); and Ausubel and Deneckere (1986). Papers on the alternating-offer bargaining game with one-sided incomplete information include: Grossman and Perry (1986); Gul and Sonnenschein (1988); and Ausubel and Deneckere (1988). Papers on one-sided incomplete information bargaining games with other extensive forms include: Rubinstein (1985); and Admati and Perry (1987). Excellent reviews of the sequential bargaining literature can be found in Rubinstein (1987) and Wilson (1987).
6. It would be literally the same incentive constraint, but for the fact that the seller's conditional distribution about the buyer's valuation may have evolved.

7. We attribute the "flavor" of the equilibrium to Chatterjee and Samuelson (1987), but we do not mean to imply that they would advocate this precise equilibrium description in the game with continuous offers and types. In particular, because this description uses two constant price paths, $\pi_1(t) = \underline{b}$ and $\pi_2(t) = \bar{s}$, this equilibrium has the property that if $\underline{b} < s < b < \bar{s}$, then the two players never trade. We imagine that Chatterjee and Samuelson might prefer a more complicated description utilizing descending price paths. See also their sequel paper, Chatterjee and Samuelson (1986).
8. In the terminology we introduce later, the Chatterjee-Samuelson (1987) equilibrium, unlike the Cramton (1984) or Cho (1988) equilibria, does not induce a seller-first mechanism.
9. The terminology here is from Cramton (1985).
10. Theorem 1 does not literally say that a "folk theorem" does not hold for seller-first mechanisms. To argue the latter, it is necessary to show that there exists a seller-first mechanism $\{p, x\}$ such that: no ICBM $\{p', x'\}$ yielding the same interim utilities as $\{p, x\}$ is implementable. However, we will demonstrate in a future paper that $\{p^4, x^4\}$ actually is a counterexample to a folk theorem for seller-first mechanisms.
11. Somewhat more compactly, one could define stationarity as follows: suppose two different histories induce the same beginning-of-period beliefs. Then the sequential equilibria induced on the remainder of the game should coincide.
12. Observe that, in the seller-offer extensive form, the stationarity restriction on the seller's updating rule is largely redundant. After any history of positive prices, the buyer's strategy must call for a positive probability of rejection. Bayes' rule then completely pins down the seller's beliefs, because of the absence of observable out-of-equilibrium buyer behavior (excluding acceptances, which end the game). The stationarity of the buyer's acceptance decision makes the seller's updating rule stationary as well.

13. Cho's equilibrium does have the property that sellers inflate their price offers not just in the initial period, but also in all subsequent periods. This separation after separation occurs because the seller must keep on convincing the buyer that she is telling the truth. It does give the equilibrium the desirable property that the number of out-of-equilibrium moves that can be made in any period is minimal.
14. There is, of course, no reason that the introductory price of seller types exceeding zero should converge to their valuations: an effort to cut prices in the hope of accelerating sales could be interpreted by the buyer as a signal of low valuation, making him reluctant to buy, and making discounts unprofitable to start with.
15. Theorem 2 uses our earlier assumption that the minimum values in the supports of $F_1(\cdot)$ and $F_2(\cdot)$ are the same (i.e., $\underline{s} = 0 = \underline{b}$). In all fairness, it should be observed that the equilibrium constructions of Cho and Cramton are also defined when $\underline{b} > 0$. Nevertheless, in that case, (9) still holds. Observe now that the solution to $\max \int_0^1 \bar{p}_1(s)g(s)ds$ subject to $\int_0^1 \bar{p}_1(s)ds = k$ occurs at $\bar{p}_1(s) = 1$, for $s \leq k$ and $\bar{p}_1(s) = 0$ for $s > k$. Since $U_1^n(0,0) = \int_0^1 \bar{p}_1^n(s)ds = k_n$, where $k_n \downarrow \underline{b}$, this establishes an upper bound for the right side of (9) equal to $G(k_n)$. Hence $\int_0^1 \bar{p}_1^n(s)d\mu_1(s) \leq G(k_n) \downarrow G(\underline{b})$, and so the mechanism is still very inefficient when \underline{b} is not too large. It is a corollary of this result that when $\underline{b} < \bar{s}$ and in any sequence of sequential equilibria in which the Coase conjecture property holds for the seller of type \underline{b} , the "no-delay" result (see Gul and Sonnenschein, 1988) cannot hold.
16. The reader may observe that in order to implement mechanisms which are essentially fully revealing (such as the monopoly mechanism), the amount of pooling will have to decrease as the time between offers approaches zero. This is indeed accomplished in our equilibrium construction.
17. It seems that we only use: there exists $\bar{\epsilon} > 0$ such that, whenever $0 < \epsilon < \bar{\epsilon}$, $\tilde{U}_1(s',s) \leq \tilde{U}_1(s + \epsilon, s)$ for every $s' \geq s + \epsilon$ and $\tilde{U}_1(s',s) \leq \tilde{U}_1(s - \epsilon, s)$ for every $s' \leq s - \epsilon$.

18. This particular equilibrium construction satisfies the heightened sense of stationarity on account that seller types "drop off" the path $\pi_0(\cdot)$ in every period, so that buyer beliefs about the seller become a sufficient statistic for the state, along the equilibrium path. Note also that this heightened stationarity is analogous to Gul, Sonnenschein and Wilson's (1986) definition for one-sided incomplete information where buyer acceptance behavior depends only on the current seller offer and not on the seller's beliefs about the buyer.
19. The reader might wonder what the most efficient seller-first mechanism is when F_1 is not convex or F_2 is not concave. While we do not have a general answer to this question, we have established that when $F_1(v_1) = v_1$ and $F_2(v_2) = 1 - (1 - v_2)^\gamma$, then for $\gamma > 1$, the optimal seller-first mechanism is the single-price mechanism that maximizes the gains from trade. It would not be too surprising if single-price mechanisms were optimal more generally. It is also worth noting that, for $\gamma = 1$, the efficiencies of the monopoly mechanism and the optimal single-price mechanism coincide. In fact, in that case there exists a continuum of other 0-1 mechanisms that yield equal efficiency (all with $p(v_1, v_2) = 0$ for $v_2 < 1/2$ and $p(0, v_2) = 1$ for $v_2 > 1/2$).
20. For example, consider a durable goods monopoly whose stock is not publicly traded and, so, whose financial statements do not become common knowledge. Or suppose that the financial statements are public but are not entirely illuminating about marginal cost.
21. But not completely. Under the distributional assumptions of Theorem 4, it may be unnecessary to regulate durable goods monopoly, as the monopoly mechanism is relatively efficient.

Appendix

Proof of Theorem 1: We will first prove the theorem when the distribution function $F_1(\cdot)$ is absolutely continuous and when $p(\cdot, \cdot)$ satisfies two additional conditions: $\bar{p}_1(\cdot)$ is a strictly decreasing function on $[0, \hat{s}]$; and for every $s \in [0, \hat{s}]$, $p(s, 1) = 1$ and $p(s, \cdot)$ is a continuous function of b on the set $[\hat{b}(s), 1]$ (where $\hat{b}(s) \equiv \inf\{b: p(s, b) > 0\}$). We make these extra assumptions in Parts I-IV of the proof; in part V, we complete the argument by relaxing these assumptions.

Part I: Construction of N price paths which replicate the mechanism.

Let $\{\epsilon^n\}_{n=1}^{\infty}$ be any sequence such that $0 < \epsilon^n < 1$, for all $n \geq 1$, and $\epsilon^n \downarrow 0$. For every n , define: $\hat{s}^n = \sup\{s \in [0, \hat{s}]: \hat{b}(s') - s' > \epsilon^n \text{ for all } s' \in [0, s]\}$. By part (iii) of Assumption 2, $\hat{s}^n \uparrow \hat{s}$.

For any ϵ^n belonging to $\{\epsilon^n\}_{n=1}^{\infty}$, select arbitrary N satisfying $N > 3/\epsilon^n$. We define a grid $\{s_k\}_{k=1}^N$ of N sellers by $s_k = (k/N)\hat{s}^n$, $k = 1, \dots, N$. We will now generally suppress the "n" from our notation until Part IV.

For every $z > 0$ and for every k ($1 \leq k \leq N$), we will define a discrete-time price path for each seller, s_k , from the grid. Let $b_k(t) \equiv \sup\{b: p(s_k, b) \leq e^{-rt}\}$ be the highest buyer valuation remaining in the market at time t when the seller is s_k . We define:

$$(A.1) \quad \pi_k^z(t) = (1 - \delta) \sum_{j=0}^{\infty} \delta^j b_k(t + (j+1)z), \text{ for } t = 0, z, 2z, \dots,$$

where $\delta \equiv e^{-rz}$. In the unlikely event that the $\pi_k^z(0)$ are not distinct for

$k = 1, \dots, N$, we raise the $b_k(z)$ arbitrarily slightly so as to redefine the $\pi_k^z(0)$ and make them distinct. Then, the seller's initial offer in (A.1) fully reveals k to the buyer.

As shown in Stokey (1981), the price path given by (A.1) induces exactly the interval $(b_k((i+1)z), b_k(iz)]$ of buyer types to purchase in period i . Furthermore, our condition that $p(s,1) = 1$ and $p(s,\cdot)$ is continuous on $[\hat{b}(s),1]$ implies that $b_k((i+1)z) < b_k(iz)$ if and only if $b_k(iz) > \hat{b}(s_k)$. Thus, there are positive sales in every period until trade concludes.

We next argue that for any $\alpha > 0$, there exists $\bar{z}(\alpha)$ such that for every z ($0 < z < \bar{z}(\alpha)$):

$$(A.2) \quad \begin{aligned} |p^z(s_k, b) - p(s_k, b)| &< \alpha, & \text{for all } k = 1, \dots, N \\ |x^z(s_k, b) - x(s_k, b)| &< \alpha, & \text{and for all } b \in [0,1], \end{aligned}$$

where $p^z(s_k, \cdot)$ and $x^z(s_k, \cdot)$ are the purchase probabilities and expected transfers implicit in (A.1). Define $\bar{z}(\alpha)$ by $1 - e^{-r\bar{z}(\alpha)} = \alpha/2$. Also, define $t_k(b)$ to be the "time" at which buyer b purchases in the mechanism $\{p, x\}$ against seller s_k , and for any z ($0 < z < \bar{z}(\alpha)$), define $t_k^z(b)$ to be the time at which buyer b purchases from the discrete price path $\pi_k^z(t)$. Observe, from (A.1), that we have defined $\pi_k^z(t)$ so that, for all $i = 0, 1, 2, \dots$, the intervals $(b_k((i+1)z), b_k(iz)]$ purchase at time iz . Consequently, $t_k^z(b) = \max\{mz: mz \leq t_k(b) \text{ and } m \text{ is an integer}\}$, for all $b \in [0,1]$. Then, $0 \leq t_k(b) - t_k^z(b) < z$ for all $b \in [0,1]$, implying $0 \leq \exp\{-rt_k^z(b)\} - \exp\{-rt_k(b)\} < 1 - e^{-rz} < 1 - e^{-r\bar{z}(\alpha)} = \alpha/2$. Thus $|p^z(s_k, b) - p(s_k, b)| < \alpha/2$. Since $bp(s_k, b) - x(s_k, b) = \int_0^b p(s_k, r)dr$ and

similarly for the discrete approximation, we have:

$$|x^z(s_k, b) - x(s_k, b)| \leq b|p^z(s_k, b) - p(s_k, b)| \\ + \int_0^b |p^z(s_k, r) - p(s_k, r)| dr \leq \alpha$$

as desired.

We further argue that inequalities (A.2) imply:

$$(A.3) \quad |\tilde{U}_1^z(s_k, s) - \tilde{U}_1(s_k, s)| < 4\alpha, \text{ for } k = 1, \dots, N \text{ and } s \in [0, 1].$$

Using (6) and (A.2), observe that:

$$(A.4) \quad |\tilde{U}_1^z(s_k, s) - \tilde{U}_1(s_k, s)| = \\ = \left| \int_{\hat{b}(s_k, s)}^1 \{ [x^z(s_k, b) - sp^z(s_k, b)] - [x(s_k, b) - sp(s_k, b)] \} d\mu_2(b) \right. \\ \left. + \int_{\hat{b}^z(s_k, s)}^{\hat{b}(s_k, s)} [x^z(s_k, b) - sp^z(s_k, b)] d\mu_2(b) \right| \\ \leq \int_{\hat{b}(s_k, s)}^1 2\alpha d\mu_2(b) + \int_{\hat{b}^z(s_k, s)}^{\hat{b}(s_k, s)} |x^z(s_k, b) - sp^z(s_k, b)| d\mu_2(b),$$

where $\hat{b}^z(\cdot)$ is defined analogously to $\hat{b}(\cdot)$. Observe that everywhere in the interval between $\hat{b}^z(s_k, s)$ and $\hat{b}(s_k, s)$, $[x^z(s_k, b) - sp^z(s_k, b)][x(s_k, b) - sp(s_k, b)] < 0$; and since the two factors are within 2α of each other, they must be within 2α of zero. In particular, $|x^z(s_k, b) - sp^z(s_k, b)| < 2\alpha$.

allowing us to conclude (A.3) from the inequalities (A.4).

Finally, define

$$(A.5) \quad \Gamma = \inf\{\tilde{U}_1(s_j, s_j) - \tilde{U}_1(s_k, s_j) : |j - k| = 1, \\ 1 \leq j \leq N \text{ and } 1 \leq k \leq N\}.$$

Observe that $\Gamma > 0$ since: (i) $\tilde{U}_1(\cdot, \cdot)$ coincides with $U_1(\cdot, \cdot)$ in (A.5), because $|s_j - s_k| < \epsilon^n$, $s_j, s_k \leq s_N$, and the definition of $s_N (\equiv \hat{s}^n)$, and (ii) $U_1(s_j, s_j) > U_1(s_k, s_j)$, because $\bar{p}_1(\cdot)$ is strictly monotone (inducing strict preference).

Let $\alpha = \Gamma/12$. Define $\bar{z}_1 \equiv \bar{z}(\alpha)$, and let z satisfy $0 < z < \bar{z}_1$. Consider any $s \in [s_k, s_{k+1}]$, where $2 \leq k \leq N - 2$. We will now demonstrate that $\arg \max_{1 \leq j \leq N} \{\tilde{U}_1^z(s_j, s)\}$ equals k or $k + 1$. First, observe that:

$$(A.6) \quad \tilde{U}_1(s_k, s) \geq \tilde{U}_1(s_j, s) + \Gamma, \text{ for all } j (1 \leq j < k),$$

$$\tilde{U}_1(s_{k+1}, s) \geq \tilde{U}_1(s_j, s) + \Gamma, \text{ for all } j (k + 1 < j \leq N).$$

This is because $\tilde{U}_1(s_i, s) = \bar{x}_1(s_i) - s\bar{p}_1(s_i)$, for $i = k - 1, k, k + 1$ and $k + 2$ (since, as in the previous paragraph, $|s_i - s| < \epsilon^n$ and so $\tilde{U}_1(\cdot, \cdot)$ coincides with $U_1(\cdot, \cdot)$). Consequently:

$$(A.7) \quad \begin{aligned} \tilde{U}_1(s_k, s) - \tilde{U}_1(s_{k-1}, s) &= [\bar{x}_1(s_k) - \bar{x}_1(s_{k-1})] \\ &+ s[\bar{p}_1(s_{k-1}) - \bar{p}_1(s_k)] \geq [\bar{x}_1(s_k) - \bar{x}_1(s_{k-1})] \\ &+ s_k[\bar{p}_1(s_{k-1}) - \bar{p}_1(s_k)] = \tilde{U}_1(s_k, s_k) - \tilde{U}_1(s_{k-1}, s_k). \end{aligned}$$

Thus, $\tilde{U}_1(s_k, s) - \tilde{U}_1(s_{k-1}, s) \geq \Gamma$. Similarly, $\tilde{U}_1(s_{k+1}, s) - \tilde{U}_1(s_{k+2}, s) \geq \Gamma$.

By Assumption 2(i), this implies (A.6). Second, from (A.3),

$|\tilde{U}_1^Z(s_k, s) - \tilde{U}_1(s_k, s)| < \Gamma/3$, for $1 \leq k \leq N$ and $s \in [0, 1]$. This together

with (A.6) immediately implies:

$$\tilde{U}_1^Z(s_k, s) > \tilde{U}_1^Z(s_j, s) + \Gamma/3, \text{ for all } j \text{ (} 1 \leq j < k \text{), and}$$

(A.8)

$$\tilde{U}_1^Z(s_{k+1}, s) > \tilde{U}_1^Z(s_j, s) + \Gamma/3, \text{ for all } j \text{ (} k + 1 < j \leq N \text{),}$$

establishing that s_k and s_{k+1} are the most preferred options (in the $\tilde{U}_1^Z(\cdot, \cdot)$ sense) for all $s \in [s_k, s_{k+1}]$, where $2 \leq k \leq N - 2$.

Also consider any $s \in [s_1, s_2]$. The second inequality of (A.8) still holds, implying that s_1 and s_2 are her most preferred options. Similarly, consider any $s \in [s_{N-1}, s_N]$. The first inequality of (A.8) still holds, implying that s_{N-1} and s_N are her most preferred options. We have thus shown that for all $s \in [s_1, s_N]$, a seller prefers the $\tilde{U}_1^Z(s_k, s)$ associated with the two nearest grid points to all other alternatives.

Part II: The treatment of the top and the bottom.

We still need to treat the two remaining intervals ($[0, s_1]$ and $[s_N, 1]$) of sellers. First, we treat the upper interval: we add extra price paths such that all seller types in the interval $(s_N, 1]$ prefer one of the extra price paths to all of the $\pi_k^Z(t)$, $1 \leq k \leq N$ (in the $\tilde{U}_1^Z(\cdot, \cdot)$ sense); and each seller type in the interval $[0, s_N)$ prefers one of the latter price paths to all of the extra price paths.

Define $\tilde{U}_1^Z(s) = \max_{1 \leq k \leq N} \tilde{U}_1^Z(s_k, s)$. We will argue that $\tilde{U}_1^Z(s_k, s)$ is a

convex function of s , for all k ($1 \leq k \leq N$), and therefore $\tilde{U}_1^Z(\cdot)$ is convex. For any $s' \in [0,1]$, define $\tilde{p}_1(s_k, s') = \int_{b(s_k, s')}^1 p(s_k, b) d\mu_2(b)$ and $\tilde{x}_1(s_k, s') = \int_{b(s_k, s')}^1 x(s_k, b) d\mu_2(b)$. Every seller $s \in [0,1]$ has the option of using the same stopping rule as s' ; consequently $\tilde{U}_1^Z(s_k, s) \geq \tilde{x}_1(s_k, s') - s\tilde{p}_1(s_k, s')$, for all s and s' . Furthermore, $\tilde{p}_1(s_k, s')$ is weakly monotone decreasing in s' . Since the family of lines $\{\tilde{x}_1(s_k, s') - s\tilde{p}_1(s_k, s')\}_{s' \in [0,1]}$ supports $\tilde{U}_1^Z(s_k, \cdot)$ from below, this establishes convexity of $\tilde{U}_1^Z(s_k, \cdot)$.

Let us define the monopoly price correspondence $g(s) = \arg \max_{\pi} [1 - F_2(\pi)][\pi - s]$ and let $M = \{\pi: \pi \in g(s) \text{ for some } s \in [0,1]\}$. Now define:

$$(A.9) \quad q^Z = \sup\{\pi \in M: [1 - F_2(\pi)][\pi - s_N] > \tilde{U}_1^Z(s_N)\}.$$

Let us denote by Π^Z the set of constant price paths $\pi(t) \equiv q$ for some $q \in [q^Z, 1]$. We will augment the existing menu of price paths by permitting the seller to offer elements of Π^Z .

First consider the case where $[1 - F_2(q^Z)][q^Z - s_N] = \tilde{U}_1^Z(s_N)$. Define $\hat{U}_1^Z(s)$ to be the highest utility attainable by type s when selecting from the menu Π^Z of price paths. Observe that there exists $s' \geq s_N$ whose monopoly price equals q^Z . Hence, s_N prefers q^Z to all higher constant price paths and so $\hat{U}_1^Z(s_N) = \tilde{U}_1^Z(s_N)$. Moreover, $\hat{U}_1^Z(\cdot)$ coincides with monopoly profits on the interval $[s', 1]$ and is linear with slope $-(1 - F_2(q^Z))$ on the interval $[0, s']$. Note that $\tilde{U}_1^Z(s') \leq \hat{U}_1^Z(s')$; since $\tilde{U}_1^Z(s_N) = \hat{U}_1^Z(s_N)$ and $\tilde{U}_1^Z(\cdot)$ is convex, this allows us to conclude that $\tilde{U}_1^Z(s) \leq \hat{U}_1^Z(s)$ for $s \in [s_N, 1]$ and $\tilde{U}_1^Z(s) \geq \hat{U}_1^Z(s)$ for $s \in [0, s_N]$. Thus, we have shown that every type

$s \in [s_1, s_N]$ selects a price path from $\{\pi_k^z\}_{k=1}^N$ and every type $s \in (s_N, 1]$ selects a price path from Π^z .

Second, consider the case where $[1 - F_2(q^z)][q^z - s_N] > \tilde{U}_1^z(s_N)$. For N sufficiently large, there exists t_N ($s_1 < t_N < s_N$) such that $[1 - F_2(q^z)][q^z - t_N] = \tilde{U}_1^z(t_N)$. As above, $\tilde{U}_1^z(s) \leq \hat{U}_1^z(s)$ for $s \in [t_N, 1]$ and $\tilde{U}_1^z(s) \geq \hat{U}_1^z(s)$ for $s \in [0, t_N]$. Observe that $\lim_{N \rightarrow \infty} t_N = \hat{s}$. Hence, we may pretend without loss of generality that we began the proof using t_N in place of s_N , and we proceed as before.

We now treat $[0, s_1]$, the bottom interval of seller types, by introducing an additional price path which begins by approximating the zero type's price path from the mechanism, but ends with an exponential rate of descent toward zero. Recall that $\hat{b}(0) = \inf\{b: p(0, b) > 0\}$ and define $\tilde{b}(0) > \hat{b}(0)$ such that:

$$(A.10) \quad \int_{\tilde{b}(0)}^1 x(0, b) d\mu_2(b) > U_1(s_1, 0).$$

This is possible because $U_1(0, 0) > U_1(s_1, 0)$ and by the absolute continuity of $F_2(\cdot)$. For every $z > 0$, define $T_0^z = \max\{nz: e^{-rnz} \geq p(0, \tilde{b}(0))\}$ and n is an integer}. T_0^z will denote the time at which we begin the use of exponentially-descending prices. For $t \leq T_0^z$ (and integer multiples of z), define $b_0(t) = \sup\{b: p(0, b) \leq e^{-rt}\}$. For $t > T_0^z$ and arbitrary $\lambda > 0$, define $b_0(t)$ by:

$$(A.11) \quad b_0(t) - \pi_0^{z, \lambda}(t - z) = e^{-rz} [b_0(t) - \pi_0^{z, \lambda}(t)],$$

where $\pi_0^{z, \lambda}(t) = e^{-\lambda(t - T_0^z)} \pi_0^{z, \lambda}(T_0^z)$, for $t \geq T_0^z$, and $\pi_0^{z, \lambda}(T_0^z)$ is uniquely

defined such that $b_0(T_0^z + z) = \tilde{b}(0)/(1 + z)$. Finally, define the entire price path $\pi_0^{z,\lambda}(t)$, $t = 0, z, 2z, \dots$, using equation (A.1).

We will now demonstrate (for small z and λ) that there exists a "cutoff seller" $c \in (0, s_1)$ such that all $s \in [0, c]$ select $\pi_0^{z,\lambda}(\cdot)$ whereas all $s \in (c, 1]$ will select one of the previously-defined price paths in $\{\pi_k^z\}_{k=1}^N \cup \Pi^z$. However, the sellers $[0, c]$ cannot pool forever, since for every $s \in (0, c]$ there will come a time when $\pi_0^{z,\lambda}(t) < s$. In the first period at which this occurs, seller s will separate from the pool and revert to making nonserious offers forever. This makes buyer optimization against $\pi_0^{z,\lambda}(\cdot)$ somewhat different from optimization against the earlier defined price paths (in which sellers pool forever). Let t be any time by which sellers have already begun to drop off. An optimizing buyer, in deciding whether to purchase at t , must take into account that the subsequent price ($\pi_0^{z,\lambda}(t + z)$) will only be offered with probability $F_1[\pi_0^{z,\lambda}(t + z)]/F_1[\pi_0^{z,\lambda}(t)]$. Thus, the buyer uses an implicit discount factor of $\delta F_1[\pi_0^{z,\lambda}(t + z)]/F_1[\pi_0^{z,\lambda}(t)]$ between times t and $t + z$. Finally, define $W(c, z, \lambda; s)$ to be the expected payoff of seller s from following the price path $\pi_0^{z,\lambda}(\cdot)$ when the buyer and seller optimize as above and the buyer believes that the cutoff seller is c .

Given any believed cutoff seller c , we will construct an actual cutoff seller, $\gamma(c)$, and show that $\gamma(\cdot)$ has a fixed point. First, observe that $\tilde{U}_1^z(\cdot)$ is linear for $s \in [0, s_1]$, while $W(c, z, \lambda; \cdot)$ is a convex continuous function of s . Thus, $\tilde{U}_1^z(s) = W(c, z, \lambda; s)$ for at most two values $s \in [0, s_1]$. Second, observe using (A.10) that there exists \bar{z}_2 ($0 < \bar{z}_2 < \bar{z}_1$) such that $W(c, z, \lambda; 0) > \tilde{U}_1^z(0)$ whenever $0 < z < \bar{z}_2$ and $0 < \lambda < \bar{z}_2$. Also, by incentive compatibility, there exists \bar{z}_3 ($0 < \bar{z}_3 < \bar{z}_2$) such that $\tilde{U}_1^z(s_1) > W(c, z, \lambda; s_1)$

whenever $0 < z < \bar{z}_3$ and $0 < \lambda < \bar{\lambda}_3$. Consequently, there exists one and only one intersection point of $\tilde{U}_1^z(\bullet)$ and $W(c, z, \lambda; \bullet)$, for such z and λ . We denote this by $\gamma(c)$. Third, note that $W(c, z, \lambda; s)$ is continuous in c , for fixed z , λ and s . This follows from the fact that interperiod discount factors, and hence buyer purchases, are continuous in the cutoff c . Fourth, the Lipschitz continuity (with constant of one) of $\tilde{U}_1^z(\bullet)$ and $W(c, z, \lambda; \bullet)$ implies that $\gamma(\bullet)$ is a continuous function. The Brouwer fixed point theorem thus establishes the existence of a value $c \in (0, s_1)$ such that $\gamma(c) = c$. For future use, let us furthermore note that there exists $\chi > 0$ such that:

$$(A.12) \quad W(c, z, \lambda; s) \geq \tilde{U}_1^z(s) + \chi, \text{ for all } s \in [0, c/2]$$

$$\text{and } (0, 0) < (z, \lambda) < (\bar{z}_3, \bar{\lambda}_3).$$

To summarize what we have shown thus far: we have defined a menu of price paths $\{\pi_k^z(t)\}_{k=1}^N \cup \pi_0^{z, \lambda}(t) \cup \Pi^z$, along with a partition $0 < s_1 < \dots < s_N < 1$, such that if all seller types are confined to initially select one of these paths and subsequently deviate only by stopping their sales, then:

- (a) Every seller type $s \in [s_k, s_{k+1}]$ for $k = 1, \dots, N - 1$ initially selects $\pi_k^z(\bullet)$ or $\pi_{k+1}^z(\bullet)$ and never subsequently deviates;
- (b) Every seller type $s \in (s_N, 1]$ initially selects a path from Π^z and never subsequently deviates; and
- (c) Every seller type $s \in [0, c]$ initially selects $\pi_0^{z, \lambda}(\bullet)$ and follows this path until such time that the price drops below her valuation. Also, every $s \in (c, s_1)$ selects $\pi_1^z(\bullet)$, and never subsequently deviates,

whenever $(0, 0) < (z, \lambda) < (\bar{z}_3, \bar{\lambda}_3)$. Furthermore, the buyer behavior used in

the above calculations of seller utility is optimal, given the prescribed seller behavior. In order to complete the proof, we need to specify behavior which deters the seller from deviating in a way other than selecting one of the prescribed price paths and subsequently stopping all sales.

Part III: Demonstration that a stationary sequential equilibrium is formed.

We now complete the description of the sequential equilibrium by specifying the beliefs and actions which are triggered by a particular type of (detectable) seller deviation. Suppose that the seller deviates from one of the prescribed price paths yet persists in attempting to make sales. Then the buyer updates his beliefs to $s = 0$ (with probability one) and maintains those beliefs forever after. Furthermore, the buyer adopts the buyer strategy from a weak-Markov equilibrium in the game of one-sided incomplete information where $s = 0$ and the distribution function of buyer valuations equals $F_2(\cdot)$. Finally, the seller maximizes against this buyer strategy.

The existence of such weak-Markov equilibria is guaranteed (for every $z > 0$) by Theorem 4.2 of Ausubel and Deneckere (1986). Let the "state" b refer to the highest remaining buyer valuation. In any weak-Markov equilibrium, we may define the "Coase price" (denoted $\pi_{\text{Coase}}(b)$) to be the price which the seller is supposed to charge, and the "choke price" (denoted $\pi_{\text{choke}}(b)$) to be the infimum of all prices which induce zero probability of sales, when the state is b . Since (by Assumption 1(iii)) $Mx^\alpha \leq F_2(x) \leq Lx^\alpha$ for all $x \in [0,1]$, the uniform version of the Coase conjecture (Theorem 5.4 of Ausubel and Deneckere, 1986) implies that, for every $\mu > 0$, there exists

$\bar{z}(\mu) > 0$ such that $\pi_{\text{Coase}}(b) < \mu b$, for all $b \in (0,1]$, whenever $0 < z < \bar{z}(\mu)$. Observe that the buyer of valuation b is indifferent between prices of $\pi_{\text{choke}}(b)$ this period and $\pi_{\text{Coase}}(b)$ next period. Consequently, there exists \bar{z}_4 ($0 < \bar{z}_4 < \bar{z}_3$) such that $\pi_{\text{choke}}(b)/b < \min\{c/2, \chi\}$, for all $b \in (0,1]$, whenever $0 < z < \bar{z}_4$ (where χ is taken from inequality (A.12)).

The previous sentence proves that, when $0 < z < \bar{z}_4$, it is impossible for any $s \in [c/2, 1]$ to generate any positive profits after deviating from an initially-selected price path, a fact which we had only been assuming until now. If $s \in [0, c/2)$ and the seller initially followed one of $\{\pi_k^z\}_{k=1}^N \cup \Pi^z$, we cannot preclude positive profits after a subsequent deviation. However, when $0 < z < \bar{z}_4$, the profits after such a deviation are bounded by $\pi_{\text{choke}}(b)$ (which is less than χ); inequality (A.12) guarantees that any $s \in [0, c/2)$ would have done better by initially choosing $\pi_0^{z, \lambda}(t)$, where $(0,0) < (z, \lambda) < (\bar{z}_4, \bar{z}_4)$.

The only deviations which remain to be excluded are of the form: some $s \in [0, c/2)$ initially follows $\pi_0^{z, \lambda}(t)$ but deviates in some subsequent period. We now demonstrate that these remaining deviations are unprofitable by calculating a lower bound on the profits to seller s from continuing to adhere to $\pi_0^{z, \lambda}(t)$ when the state is b , and showing that it exceeds "Coase profits" for all s and all b , when z is sufficiently small. We will argue this as if the price path $\pi_0^{z, \lambda}(\cdot)$ were everywhere exponentially descending. At time $t \geq T_0^z$, the continuation path literally has this property; a simple modification extends this argument to earlier periods.

Let π ($0 < \pi < 1$) be any initial price and let λ be the real-time rate of descent in price. If all $s \in [0, c]$ initially follow this price path and if seller s drops off in the first period that price falls below s , every

buyer $b \in (0,1]$ selects a time t to purchase by solving:

$$(A.13) \quad \max_t \{F_1[\min(c, \pi e^{-\lambda t})] e^{-rt} (b - \pi e^{-\lambda t})\}, \text{ where } t = 0, z, 2z, \dots$$

For every b , define $\underline{t}(b)$ by $\pi e^{-\lambda \underline{t}(b)} = b$ and $\bar{t}(b) = \arg \max_t \{e^{-rt} (b - \pi e^{-\lambda t})\}$. Observe, in (A.13), that b purchases no earlier than time $\max\{0, \underline{t}(b)\}$ and no later than time $\max\{0, \bar{t}(b) + z\}$. Suppose now that b_0 is the current state. Then the current time is at least $\max\{0, \underline{t}(b_0)\}$. Define $b_1 = (M/2L)^{1/\alpha} b_0$, where L , M and α are from Assumption 1(iii), and observe that the time at which b_1 purchases is no later than $\max\{0, \bar{t}(b_1) + z\}$. Consequently, the seller discounts sales to the interval $[b_1, b_0)$ using a factor no less than $e^{-r[\bar{t}(b_1) - \underline{t}(b_0) + z]} = w \delta (M/2L)^{r/\alpha \lambda}$, where $w \equiv [r/(r + \lambda)]^{r/\lambda}$. The probability that b is contained in $[b_1, b_0)$ is bounded below by $F(b_0) - F(b_1) \geq Mb_0 - Lb_1 = (M/2)b_0^\alpha$. Meanwhile, all sales to the interval $[b_1, b_0)$ occur at prices no less than $\pi e^{-\lambda \max\{0, \bar{t}(b_1) + z\}} = \min\{\pi, [r/(r + \lambda)] b_1 e^{-\lambda z}\}$. Note there exists $c_1 > 0$ such that $c_1 b_0 < \min\{\pi, [r/(r + \lambda)] (M/2L)^{1/\alpha} b_0 e^{-\lambda z}\}$ whenever $0 < z < \bar{z}_4$. Hence, a lower bound on the expected profit to s when the state is b_0 is given by:

$$(A.14) \quad V(s, b_0) = w e^{-r\bar{z}_4} (M/2L)^{r/\alpha \lambda} (M/2) b_0^\alpha [c_1 b_0 - s] \equiv c_2 b_0^\alpha (c_1 b_0 - s).$$

Let \bar{z}_5 ($0 < \bar{z}_5 < \bar{z}_4$) be sufficiently small that $\pi_{\text{choke}}(b_0) < (c_1/2)b_0$ and $\pi_{\text{Coase}}(b_0) < (c_1 c_2/2L)b_0$, for all $b_0 \in (0,1]$, whenever $0 < z < \bar{z}_5$. First, consider any (s, b_0) such that $s \geq \pi_{\text{choke}}(b_0)$. Seller s earns nonnegative profits from remaining on price path $\pi e^{-\lambda t}$ (with the usual stopping rule). Deviation triggers Coase expectations, nets zero profits to seller s , and is

hence deterred. Second, suppose $s < \pi_{\text{choke}}(b_0)$. Since $s < (c_1/2)b_0$ and by (A.14), $V(s, b_0) > (c_1 c_2 / 2) b_0^{1+\alpha}$. Profits from deviation are bounded above by $\pi_{\text{Coase}}(b_0) F_2(b_0) < (c_1 c_2 / 2) b_0^{1+\alpha}$, so deviation is again deterred, demonstrating that we have indeed constructed a sequential equilibrium.

Furthermore, the constructed equilibrium is stationary. As noted earlier, each price path $\{\pi_k^z(\cdot)\}_{k=1}^N$ induces positive sales in every period until trade concludes. Elements of Π^z are constant price paths, surely lending stationarity. Along $\pi_0^{z, \lambda}(\cdot)$, there are sales in every period in which price is no less than c ; and buyer beliefs about the seller evolve in every period that price is less than c (as sellers "drop off" every period). If the seller has deviated undetectably, her strategy prescribes that she follow her adopted price path until it drops below her valuation (or that she detectably deviate); the same argument as above shows stationarity. If the seller has deviated detectably, this fact is reflected in the state, as only histories with detectable deviations lead to buyer beliefs that $s = 0$. Strategies after a detectable deviation are certainly stationary, as buyers merely use reservation price strategies and sellers optimize against them.

Part IV: Completion of the proof, subject to the extra assumptions.

We have shown that, for every $n \geq 1$, there exists \bar{z}^n ($\equiv \bar{z}_5^n$ of Part III, for $\epsilon = \epsilon^n$) such that for all time intervals z between offers ($0 < z < \bar{z}^n$) and all rates of descent λ for the bottom pool ($0 < \lambda < \bar{z}^n$), the above construction yields a sequential equilibrium satisfying inequalities (A.2) with $\alpha = \Gamma/12$ and Γ defined by (A.5). Observe that for $k \geq j$, $\tilde{U}_1(s_j, s_j) \geq \tilde{U}_1(s_k, s_j) \geq \tilde{U}_1(s_k, s_k) \geq \tilde{U}_1(s_j, s_j) - (s_k - s_j)$, where the first inequality follows from incentive compatibility, the second follows from the fact that

$s_j \leq s_k$, and the third follows from the fact that $U_1(s) = \int_s^1 \bar{p}_1(t) dt$. Also observe that for $\lambda \leq j$, $\tilde{U}_1(s_j, s_j) \geq \tilde{U}_1(s_\lambda, s_j) \geq \tilde{U}_1(s_\lambda, s_\lambda) - (s_j - s_\lambda) \geq \tilde{U}_1(s_j, s_j) - (s_j - s_\lambda)$, where the first inequality follows from incentive compatibility, the second follows from the fact that $\tilde{U}_1(s_\lambda, s_j) \geq \bar{x}_1(s_\lambda) - s_j \bar{p}_1(s_\lambda)$, and the third follows from the monotonicity of $U_1(\cdot)$. Hence, $|\tilde{U}_1(s_j, s_j) - \tilde{U}_1(s_k, s_j)| < |s_j - s_k|$ for all j, k ; since $|s_j - s_k| < \epsilon^n/3$ whenever $|k - j| = 1$, we conclude that $\Gamma < \epsilon^n/3$ and so $\alpha < \epsilon^n/36$. We select a sequence of pairs $\{z^n, \lambda^n\}_{n=1}^\infty$ such that $(0, 0) < (z^n, \lambda^n) < (\bar{z}^n, \bar{z}^n)$ for all $n \geq 1$ and $(z^n, \lambda^n) \downarrow (0, 0)$, and we denote by $\{\sigma^n\}_{n=1}^\infty$ the associated sequence of stationary sequential equilibria. Also let $\{p^n, x^n\}$ be the ICBM induced by σ^n . We will now show that $\{\sigma^n, z^n\}_{n=1}^\infty$ implements the ICBM $\{p, x\}$.

In any equilibrium σ as constructed above, let $\bar{b}^n \equiv q^{z^n}$ denote the value defined in (A.9), for $\{\sigma^n, z^n\}$. Let $\lambda(\cdot)$ denote (one-dimensional) Lebesgue measure. Define $A^n = \{(s, b) : s_1 \leq s \leq s_N \text{ and } 0 \leq b \leq 1\}$ and $A'^n = \{(s, b) : \hat{s} \leq s \leq 1 \text{ and } 0 \leq b \leq \bar{b}^n\}$. We will now show that:

$$(A.15) \quad \lambda\{s \in [s_1, s_N] : |p^n(s, b) - p(s, b)| > 2\sqrt{\epsilon^n}\} < (K/3)\sqrt{\epsilon^n},$$

for all $b \in [0, 1]$.

Since $p^n(s, b) = 0 = p(s, b)$ on A'^n , since the $(\lambda \times \lambda)$ -measure of the complement of $(A^n \cup A'^n)$ converges to zero, and since $\epsilon^n \downarrow 0$, this will establish that p^n converges to p in $(\lambda \times \lambda)$ -measure on the unit square. Furthermore, since $F_1(\cdot)$ and $F_2(\cdot)$ are absolutely continuous distribution functions, this also establishes that $\{\sigma^n, z^n\}_{n=1}^\infty$ implements the ICBM $\{p, x\}$.

Define $J^n(b) \equiv \{j : 1 \leq j \leq N - 1 \text{ and the total variation of } p(\cdot, b) \text{ on } [s_j, s_{j+1}] \text{ exceeds } \sqrt{\epsilon^n}\}$, for every $n \geq 1$ and $b \in [0, 1]$. Observe that the

cardinality of $J^n(b)$ is less than $K/\sqrt{\epsilon^n}$; otherwise, the total variation of $p(\cdot, b)$ on $[s_1, s_N]$ would exceed K , violating Assumption 2(ii). For fixed n and b , consider any interval $[s_j, s_{j+1}]$ such that $1 \leq j \leq N - 1$ and $j \notin J^n(b)$. Since for any $s \in [s_j, s_{j+1}]$, $p^n(s, b)$ equals $p^n(s_j, b)$ or $p^n(s_{j+1}, b)$, we have (for either $i = j$ or $i = j + 1$):

$$\begin{aligned} |p^n(s, b) - p(s, b)| &= |p^n(s_i, b) - p(s, b)| \\ &\leq |p^n(s_i, b) - p(s_i, b)| + |p(s_i, b) - p(s, b)| \\ &\leq \epsilon^n/36 + \sqrt{\epsilon^n} < 2\sqrt{\epsilon^n}. \end{aligned}$$

Thus, $s \in [s_j, s_{j+1}]$ can only be an element of the set defined in (A.15) if $j \in J^n(b)$. But $\#\{s: s \in [s_j, s_{j+1}] \text{ and } j \in J^n(b)\} \leq (K/3)\sqrt{\epsilon^n}$, since $s_{j+1} - s_j < \epsilon^n/3$, for all j ($1 \leq j \leq N - 1$), proving (A.15). This completes the proof of the theorem under the assumptions of the first sentence of the proof.

Part V: Completion of the proof, without the extra assumptions.

First let us relax the assumption that $p(s, 1) = 1$ and $p(s, \cdot)$ is continuous on $[\hat{b}(s), 1]$, for all s . Observe that the previous construction still yields sequential equilibria $\{\sigma^n\}_{n=1}^\infty$ and induced ICBM's $\{p^n, x^n\}_{n=1}^\infty$; however σ^n are not necessarily stationary. We now argue that each $\{p^n, x^n\}$ is itself implementable by stationary sequential equilibria; a diagonal argument then clearly demonstrates the implementability of $\{p, x\}$ in stationary equilibria. Recall that each $\{p^n, x^n\}$ considers some grid s_1, \dots, s_N of sellers. Let $\theta^i \downarrow 0$: observe that for every $n \geq 1$ and $i \geq 1$,

we can construct an N -tuple $\{p^{ni}(s_k, b), x^{ni}(s_k, b)\}_{k=1}^N$ of pairs of functions of b with the property that for all k ($1 \leq k \leq N$): $p^{ni}(s_k, 1) = 1$; $p^{ni}(s_k, \cdot)$ is monotone increasing and continuous on $[\hat{b}(s_k), 1]$; $p^{ni}(s_k, b) = 0$ whenever $0 \leq b < \hat{b}(s_k)$; $x^{ni}(s_k, b) = bp^{ni}(s_k, b) - \int_0^b p^{ni}(s_k, r) dr$; and:

$$(A.16) \quad \lambda\{b \in [0, 1]: |p^{ni}(s_k, b) - p^n(s_k, b)| > \theta^i\} < \theta^i.$$

For $z > 0$, we can define (analogously to (A.1)) discrete price paths from $\{p^{ni}(s_k, \cdot), x^{ni}(s_k, \cdot)\}_{k=1}^N$ for each s_k ($1 \leq k \leq N$) which now yield positive sales in every period until trade concludes. Mimicking arguments which appeared in the early part of the proof, it is straightforward to show, for every $n \geq 1$ and every $\theta > 0$, that there exists $i^n \geq 1$ and $z^n > 0$ such that whenever $i \geq i^n$ and $0 < z < z^n$, all but measure θ of sellers sort themselves in the same way against $\{p^{ni}, x^{ni}\}$ as against $\{p^n, x^n\}$. This establishes implementability of $\{p^n, x^n\}$.

Next, let us relax the assumption that $\bar{p}_1(\cdot)$ is strictly monotone on $[0, \hat{s}]$. Construct the grid $\{s_k\}_{k=1}^N$ as before. Define β to satisfy $0 < 1 - F_2(\beta) < U_1(s_N)/3$. Let $g(s) \equiv \arg \max_{\pi} [1 - F_2(\pi)][\pi - s]$ be the monopoly price correspondence. Select γ_1 and γ_2 such that $\beta < \gamma_1 < \gamma_2 < 1$ and $g(s)$ is single-valued for all $s \in [\gamma_1, \gamma_2]$. Observe that $[g(\gamma_1), g(\gamma_2)] \subset \text{Range } g$. Define $q(\cdot)$ to be any strictly increasing C^1 function on $[0, \gamma_2]$ such that $q(0) = g(\gamma_1)$, $q(\gamma_2) = g(\gamma_2)$, and $q(s) > \sup g(s)$ in the interior of its domain. Extend $q(\cdot)$ to all of the domain $[0, 1]$ by defining $q(s) = \sup g(s)$ for $s \in (\gamma_2, 1]$. In order to make it incentive-compatible for s to name the single price $q(s)$ forever, we instruct s to delay offering that price until such time she discounts to

$h(s)$, where:

$$(A.17) \quad h(s) = \exp\{-q'(s) \int_s^{\gamma_2} \left[\frac{f_2(q(t))}{1 - F_2(q(t))} - \frac{1}{q(t) - t} \right] dt\},$$

for $0 \leq s \leq \gamma_2$, and $h(s) = 1$ for $\gamma_2 < s \leq 1$. Observe that $q(t)$ exceeds $\sup g(t)$ and $q(t) = g(t')$ for some $t' > t$. Hence,

$(\partial/\partial\pi)\{[1 - F_2(\pi)][\pi - t]\} < 0$ when evaluated at $\pi = q(t)$ (whenever $f_2(\cdot)$ is defined). It follows that the integrand of (A.17) is positive, implying that $h(\cdot)$ is strictly increasing on $[0, \gamma_2]$. Finally we define an auxiliary mechanism $\{p^A, x^A\}$ by:

$$(A.18) \quad p^A(s, b) = \begin{cases} h(s), & \text{if } b \geq q(s) \\ 0, & \text{otherwise.} \end{cases}$$

and $x^A(s, b) = p^A(s, b)q(s)$. Observe that $\{p^A, x^A\}$ is a seller-first mechanism and $\bar{p}_1^A(\cdot)$ is strictly monotone on $[0, 1]$.

Define a new (seller-first) mechanism from the original and auxiliary mechanisms by:

$$(A.19) \quad p^C(s, b) = (1 - \epsilon^n)p(s, b) + \epsilon^n p^A(s, b), \text{ and}$$

$$x^C(s, b) = (1 - \epsilon^n)x(s, b) + \epsilon^n x^A(s, b),$$

where we restrict attention to n sufficiently large that $\epsilon^n < 1/3$. Observe that $\bar{p}_1^C(\cdot)$ is strictly monotone on $[0, 1]$; construct price paths

$\{\pi_k^z\}_{k=1}^N \cup \pi_0^{z,\lambda} \cup \Pi^z$ as before, but now with reference to $U_1^C(\cdot)$. We need to argue that there exists $\bar{z} > 0$ such that sellers separate correctly whenever $0 < z < \bar{z}$. First, we claim that $\tilde{U}_1^C(s',s) - U_1^C(s',s) = (1 - \epsilon^n)[\tilde{U}_1(s',s) - U_1(s',s)]$ for all s',s such that $\tilde{U}_1^C(s',s) > 1 - F_2(\beta)$. To see this, note that seller type s (when following the price path of s' derived from $\{p^c, x^c\}$) must sell to a buyer type $b \in [0, \beta)$ to attain profits exceeding $1 - F_2(\beta)$. Since $q(s') > \beta$, s sells to the same buyer types in $[0, \beta)$ at the same prices when following the price path of s' derived from $\{p, x\}$, but accelerated by a discount factor of $1/(1 - \epsilon^n)$. Second, we claim that $\tilde{U}_1^C(s',s)$ is strictly quasiconcave in s' when $\tilde{U}_1^C(s',s) > 1 - F_2(\beta)$. This follows from: $\tilde{U}_1^C(s',s) > U_1^C(s',s) + (1 - \epsilon_n)[\tilde{U}_1(s',s) - U_1(s',s)] = \epsilon^n U_1^A(s',s) + (1 - \epsilon^n)\tilde{U}_1(s',s)$; and the fact that each of the two latter functions is strictly quasiconcave, peaking at s . Now suppose $s \in [s_k, s_{k+1}]$ where $1 \leq k \leq N - 1$. Observe that $\tilde{U}_1(s_{k+1},s) > 3[1 - F_2(\beta)]$, and since $\epsilon^n < 1/3$, $\tilde{U}_1^C(s_{k+1},s) \geq U_1^C(s_{k+1},s) > 2[1 - F_2(\beta)]$. The quasiconcavity claim demonstrates that $\tilde{U}_1^C(s_j,s) \leq \max\{\tilde{U}_1^C(s_k,s), \tilde{U}_1^C(s_{k+1},s)\}$ for all j such that $\tilde{U}_1^C(s_j,s) > 1 - F_2(\beta)$; clearly this preference still holds when $\tilde{U}_1^C(s_j,s) \leq 1 - F_2(\beta)$. Finally, for \bar{z}^n sufficiently small, the same preference holds for discrete approximations.

For each n , we can construct (by arguments analogous to those above) a stationary sequential equilibrium σ^n for some $z^n \in (0, \bar{z}^n)$. As $n \uparrow \infty$, note that the β we defined converges to one, and so the sequence $\{\sigma^n\}_{n=1}^\infty$ implements the mechanism.

Finally, let us relax the assumption that $F_1(\cdot)$ is absolutely continuous. First we consider the case where there is no mass point in the distribution at \hat{s} . The generalization requires a change in the initial specification of the grid of seller types. For every n , define $\{s_k\}_{k=1}^N$ as

at the beginning of the proof. Define $\{t_k\}_{k=1}^N$ by:

$$t_k = \inf\{t: F_1(t) \geq [k/N]F_1(\hat{s}^n)\}.$$

Now define $\{s'_k\}_{k=1}^{N'}$ to be the ordered union of $\{s_k\}_{k=1}^N$ and $\{t_k\}_{k=1}^N$, where any redundancies are omitted (i.e., $N' < 2N$). Next, repeat the main construction of the proof, only using $\{s'_k\}_{k=1}^{N'}$ in place of $\{s_k\}_{k=1}^N$. As before, we obtain a sequence of sequential equilibria with associated $\{p^n(\cdot, \cdot)\}_{n=1}^\infty$. We will now establish the analogue to inequality (A.15), where $\lambda(\cdot)$ is replaced by $\mu_1(\cdot)$.

Define $J^n(b) \equiv \{j: 1 \leq j \leq N' - 1 \text{ and the total variation of } p(\cdot, b) \text{ on } [s'_j, s'_{j+1}] \text{ exceeds } \sqrt{\epsilon^n}\}$, for every $n \geq 1$ and $b \in [0, 1]$. Still, the cardinality of $J^n(b)$ is less than $K/\sqrt{\epsilon^n}$. Observe that, for fixed n and b , $|p^n(s, b) - p(s, b)| \geq 2\sqrt{\epsilon^n}$ only if: either $s \in [0, s'_1) \cup (s'_{N'}, 1]$; or $s \in (s'_j, s'_{j+1})$ for some $j \in J^n(b)$. But $\mu_1((s'_j, s'_{j+1})) < \epsilon^n/3$ (for all $j = 1, \dots, N' - 1$) by construction, and hence $\mu_1\{s: s \in (s'_j, s'_{j+1}) \text{ and } j \in J^n(b)\} \leq (K/3)\sqrt{\epsilon^n}$. Since $F_1(\cdot)$ does not possess a mass point at 0 or \hat{s} , this establishes the implementability of $\{p, x\}$.

Second, we consider the case where there is a mass point at \hat{s} . Without loss of generality, we may assume $\hat{s} < 1$ (since in any seller-first mechanism and in any equilibrium, the probability $\bar{p}_1(1)$ that a seller of type 1 trades equals zero). Observe by the definition of \hat{s} and by $\tilde{U}_1(s, s) \geq \tilde{U}_1(s', s)$ that the price path associated with \hat{s} is necessarily a single price of \hat{s} beginning at some time T . For every $\epsilon > 0$, we modify the original mechanism by offering all seller types an additional alternative: a single price of $\hat{s} + \epsilon^2$ at a time T' , where T' is defined by $e^{-rT'} = e^{-rT}/(1 + \epsilon)$. Observe

that seller type $\hat{s} - \epsilon$ will forego this additional price path, since it would yield her utility of $[e^{-rT'}][1 - F_2(\hat{s} + \epsilon^2)][\epsilon + \epsilon^2] < [e^{-rT}][1 - F_2(\hat{s})]\epsilon = U_1(\hat{s}, \hat{s} - \epsilon)$, whereas $U_1(\hat{s} - \epsilon, \hat{s} - \epsilon) \geq U_1(\hat{s}, \hat{s} - \epsilon)$ can be obtained from the original mechanism. By the single-crossing property, no seller type $s \leq \hat{s} - \epsilon$ will select this additional price path. Hence, as $\epsilon \downarrow 0$, the ϵ -modified mechanisms converge in μ_1 -measure to the original mechanism (since \hat{s} follows a price path arbitrarily close to her original price path). If each ϵ is chosen so that there is no mass point in $F_1(\cdot)$ at $\hat{s} + \epsilon^2$, then each ϵ -modified mechanism is implementable, and a diagonal argument completes the proof.