

Efficient Actions in a Dynamic Auction Environment*

Michael Schwarz[†] and Konstantin Sonin[‡]

First draft: January, 2001

This draft: October, 2005

Abstract

We model an environment, where bidders' private values may change over time as a result of both costly private actions and exogenous shocks. Examples of private actions include investment and entry decisions; shocks might be due to exogenous changes in a potential buyer's circumstances. We describe an efficient auction mechanism that maximizes the final value of the object to the winning bidder net of the total cost of investment by all agents. In particular, we show that, assuming that the auctioneer does not have full commitment power, costly signalling is necessary for efficient entry when agents receive private information both before and after they make the entry decision. To rule out pooling equilibria that coexist with the efficient equilibrium in the basic mechanism, we introduce a virtual-implementation-style mechanism that (i) is almost efficient; (ii) forces players to coordinate on the separating equilibrium; and (iii) is simple enough to be potentially useful in practice.

JEL: D80, D44.

Keywords: auctions, efficient mechanism design.

*An earlier version of this paper circulated under the title "The Variable Value Environment: Auctions and Actions". The authors are grateful to Rick Ericson, Drew Fudenberg, Eric Maskin, Paul Milgrom, Anna Mikoucheva, Eric Rasmusen, Mike Riordan, Ilya Segal, Thomas Sjostrom, Eric Stout, Jeroen Swinkels, Robert Townsend, seminar participants at Harvard/MIT, Columbia, and participants of the 2001 NASM of the Econometric Society, the 2002 SED conference, the 2003 EEA Congress in Lausanne, and the Second Game Theory Congress for various helpful comments.

[†]University of California at Berkeley and NBER; e-mail: mschwarz@berkeley.edu.

[‡]New Economic School/CEFIR, Moscow; e-mail: ksonin@nes.ru.

1 Introduction

If an object is available for use at some future date, what is an efficient mechanism for selling it? Consider a sale of a hypothetical military base that is scheduled to close in ten years. Waiting with the sale for ten years and auctioning off the base immediately before it becomes available creates an inefficiency as the winning bidder might have missed opportunities to invest in assets complementary to the base ownership. In other words, *private actions* that a bidder chooses prior to the actual sale may influence her private value. If a bidder thinks that her likelihood of winning the object is sufficiently low, she would choose not to make a costly investment that increases her valuation of the object. On the other extreme, selling the base ten years before it becomes available seems absurd, because the expected value of the object for each bidder is likely to change over time. In other words, as long as there are privately observed *exogenous shocks* to private values that are revealed over time, auctioning off the object well in advance (before the exogenous shocks are observed) may be inefficient.

Changes in private values due to private actions and exogenous shocks are ubiquitous.¹ A sale of almost any object or service available for use at some known future date is an example of such situation, ranging from leasing a building under construction to renting a dance club for New Year’s Eve. A sale of an object in a market where search is important inevitably has elements of this environment. Consider the sale of a house that is scheduled to be auctioned off in two weeks. Before the house becomes available, bidders may have opportunities to buy other houses, essentially removing themselves from the market. Thus, we can consider an action consisting of “not buying some other reasonably priced house” as an action that boosts the value of the house. A standard auction (or any other single-round auction mechanism) is bound to be inefficient in a market with search because in this environment information revelation is necessary for achieving efficiency. Our model seeks to offer insights into auction design for such markets.

Auctions with entry costs are related to the environment we study. Mathematically, decision to pay for costly entry into an auction is a particular case of an action that boosts the private value of the agent (see Example 3). Entry costs are investigated in various contexts by Milgrom (1981), Samuelson (1985), McAfee and McMillan (1987), McAfee, Qian, and Vincent (1992), Riordan and Sappington (1987), Levin and Smith (1994), Fullerton and McAfee (1999), and Lixin Ye (2000). Both in McAfee and McMillan (1987) and Levin and Smith (1994), entry decision are made by symmetric agents that have no specific information on the value of object to them; therefore, the crucial role is played by the equilibrium number of entering agents. In contrast, Samuelson

¹Recent studies of similar environments include Bergemann and Valimaki (2000, 2002), Calzolari and Pavan (2002), and Eso and Szentes (2003). Myerson (1986) employs the mechanism design perspective to explore an environment, where only exogenous changes to agents’ valuations are possible. Other studies that focus on exogenous shocks to valuations are Haile (2000, 2001, 2003). Haile (2001) allows for the possibility of resale. McAfee, Takacs, and Vincent (1999) use a similar model.

(1985) studies an entry model where agents learn no new information upon entry; still, agents entry strategies are conditional on other agents' entry decision. Our model, where agents receive information both before and after they act (e.g., enter or make a complimentary investment) allows to extend the efficient-entry results and, at the same time, to highlight the fact that efficient entry in a general setup necessitates the use of costly signalling.

We propose a formal model with three periods. In the first period, each party receives a private signal s about its private value for the object. In the second period, a party can take a private, unobservable action (investment) at cost c that increases the value of the object by b . In the third period, bidders receive independent exogenous shocks t to their private values of the object. For i th bidder, the final reservation price for the object is a function of signals (s_i and t_i) regarding the value obtained in the first and third periods plus the benefit b from taking an action if the bidder took the action ($V(s_i, t_i) + b$) (c is sunk). Otherwise, the private value of the object is $V(s_i, t_i)$. The identity of the bidder with the highest value can only be established after the third period signals t_i are observed. Thus, an efficient allocation mechanism requires that the ownership of the object is assigned after the third period. Conducting a second price auction after the third period is not efficient (in a world where first period signals are privately observed), since it forces agents to take decisions regarding second period actions in ignorance of the expected private values of other agents. The following example illustrates this simple but essential point.

Example 1 *Let the number of participants be $N = 2$, assume that s_i are privately observed signals independently drawn from the uniform distribution on $[0, 1]$ and $V(s) = s$ (effectively there is no third period signal). Since there is no third-period uncertainty the social planner chooses exactly one agent to make a costly investment – the one with the highest first-period signal.² In a symmetric equilibrium with no revelation of the first-period signals, each agent invests if her probability of winning conditional on her own type is higher than $\frac{c}{b}$. That is, agent i invests if $s_i \geq s^* = \frac{c}{b}$. If $s^* = \frac{3}{4}$, then with probability $\frac{1}{16}$ both agents invest (which is inefficient), and with probability $\frac{9}{16}$ no agent invests (which is inefficient as well). Therefore, on average there are too few actions ($\frac{1}{2}$ instead of 1). If $s^* = \frac{1}{4}$, the situation is reversed: with probability $\frac{9}{16}$ both agents invest, and with probability $\frac{1}{16}$ no agent invests. On average, there are too many actions ($\frac{3}{2}$ instead of 1). This is hardly surprising: without signalling, there are too few investment, when actions are relatively costly ($\frac{c}{b} = \frac{3}{4}$), and there are too many investment, when actions are relatively cheap ($\frac{c}{b} = \frac{1}{4}$).*

We consider a problem of designing an efficient mechanism for allocating the object in the environment, where signals about bidder's private values (s_i , and t_i) are private. Without assuming

²In the special case of all t_i 's equal to zero, an efficient allocation rule can be implemented by assigning the ownership of the object by conducting a Vickrey auction at the end of the first period after s_i 's are privately learned. For any non-degenerate distribution of third-period signals, assigning the ownership of the object at the end of the first period is no longer efficient. Of course, the inefficiency of allocating the object at the end of the third period demonstrated by the example does not go away when t_i 's are not equal to zero.

the full commitment power on behalf of the social planner, a Vickrey-Clarke-Groves mechanism cannot be straightforwardly employed since bidders do not know their final valuations which are evolving over time and bidders can take value-enhancing actions. In our analysis, we do not assume that the auctioneer has full commitment power. If such commitment is assumed, then the dynamic modification of the VCG mechanism suggested by Bergemann and Valimaki (2001) assures a socially optimal allocation.³ Either with or without full commitment, given two-stage nature of information, efficiency requires at least two rounds of announcements of types: the announcements cannot be made at the same time, since an efficient investment decision necessarily depends on the realization of the preliminary signals.

First, for our environment, we define an outcome as the identity of the bidder who receives the object and the list of private investments taken by bidders. For a given outcome, the social surplus is equal to the value of the object to the agent who receives it net of the total cost of investment. We start with considering a social planner that pursues a strategy maximizing the expected social surplus as an efficiency benchmark. Since the exogenous shock of the third period is not known in the second period, when decisions to take actions are made, it may be efficient to have more than one bidder making an investment or to have no bidders investing at all. Part (i) of Theorem 1 establishes that if the social planner orders an agent with the first-period value s_i to invest, then she also orders all agents with value greater than s_i to invest.

Of course, an all-knowing and well-intentioned social planner is rarely available in the real world. What happens if there is no social planner but all the information is common knowledge, i.e. signals obtained by a bidder about her private value are observed by all players? Part (ii) of Theorem 1 establishes that the efficient allocation can be achieved in a decentralized case, extending results of McAfee and McMillan (1987) and Levin and Smith (1994) to the situation where potential entrants are heterogeneous and absorb information both before and after they make the entry decision.

The above results rely on bidders' private values being common knowledge. A more realistic case, where bidders privately observe their valuations, is of primary interest. Can an efficient allocation be achieved in that case? It is straightforward that an efficient allocation can not be attained without revelation of bidders' private signals (s_i) prior to the second period. If the object is allocated to the bidder with the highest value following the third period (using, say, a second price sealed bid auction) adding a cheap talk stage following the first period will not result in any information revelation and thus would lead to an inefficient outcome.⁴ Indeed, each bidder would

³The proof is spelled out in the working paper version (Bergemann and Valimaki, 2000, pages 15-17), but not in the final publication (Bergemann and Valimaki, 2002). It is also worth pointing out that the adaptation of the VCG for the dynamic environment proposed in Bergemann and Valimaki, 2000 yields large negative revenues for the seller in some states of the world, the mechanism proposed in this paper insures that the seller's revenues are positive.

⁴The condition that the object is allocated to the bidder with the highest value following the third period is a necessary, but not sufficient condition for efficient allocation. This is because efficiency depends on the set of players that invest in the second period. As we mentioned before, the winning bidder might have forgone investment

claim to be of the ‘high type’ because the higher is the perceived type of a bidder, the less likely are the other bidders to make investments and thus the lower are their subsequent bids. Theorem 2 and Theorem 3 show that there exists an efficient mechanism, where private information is revealed in the first round and the object is assigned in the second round. The first round takes place after private signals s_i are received by agents. In the first round, bidders reveal their private signals s_i by making payments (we show that the higher is the private signal s_i , the higher is the agent’s willingness to pay for reporting to other agents that the value of her private signal s_i is high). The second round consists of a second-price sealed-bid auction conducted after signals t_i are received.

As long as private signals s_i are truthfully revealed in the first round, the subgame corresponding to the second round is identical to the complete information game. Theorem 2 establishes that the mechanism described above has an efficient separating equilibrium. Unfortunately, this mechanism also has an inefficient pooling equilibrium. To rule out the pooling equilibrium, we propose a class of mechanisms that force players to coordinate on the separating equilibrium. We refer to mechanisms from this class as “ ε -coupon mechanisms.” We prove that one can always choose an ε -coupon mechanism which yields an efficient allocation with probability arbitrarily close to one. An ε -coupon mechanism consists of two rounds. The first round takes place after private signals s_i are received by agents: a non-transferable discount for amount ε is sold via a sealed-bid all-pay auction. After the all-pay auction all bids are made public. The ε discount can only be used in the second round auction. In the second round the object is sold using a Vickrey auction (if the winner of the Vickrey auction is a holder of the ε discount, she pays the second highest bid minus ε). For $\varepsilon = 0$, this mechanism is identical to the efficient mechanism described above. Theorem 4 shows that an arbitrarily small positive ε forces agents to coordinate on a separating equilibrium that yields an efficient allocation with probability converging to one as ε converges to zero. In spirit, this mechanism is very close to virtual implementation (e.g., Maskin and Sjöström, 1999).

The rest of the paper is organized as follows. In Section 2, we introduce the formal model. Section 3 describes an efficient mechanism that has a fully separating perfect Bayesian equilibrium. Section 4 introduces the ε -coupon mechanism and establishes that it has a unique robust equilibrium. Section 5 concludes.

2 The Setup

There are N *ex-ante* identical agents. In the first and the third periods, agents receive independent signals about their private values of the object.⁵ In the second period, each agent has an opportunity

opportunities enhancing the value of the object.

⁵We focus on the simplest possible model that allows to explore dynamic aspects of the environment. Extending the results to the common-value case, e.g. using the Dasgupta and Maskin (2000) constrained-efficient mechanism, is a topic for future research.

to make an irreversible costly investment that increases her private value of the object.

Timing

Period 1. Each agent receives a signal $s_i \geq 0$ about her private values, drawn independently from the same atomless distribution.

Period 2. Each agent i has an opportunity to make a (private) investment $a_i \in \{0, 1\}$, which increases the agent's private value by ba_i and costs $ca_i \geq 0$, $b > c$. When $a_i = 1$ we say that the agent i invests; if $a_i = 0$ we say that the agent i abstains from investing.

Period 3. Agents receive independent signals $t_i \geq 0$ about their private values, drawn from the same (atomless) distribution. We assume that a higher first-period signal s_i makes a higher third-period t_i (weakly) more likely. Formally, if $s_i > s'_i$, then the distribution of t_i conditional on s_i stochastically dominates the distribution of t_i conditional on s'_i (t_i and s_i being independent is a particular case).

Agent's i private value of the object equals $V_i = V(s_i, t_i)$ plus the benefit from investing. Therefore, the utility of the agent is given by:

$$U_i = \begin{cases} V(s_i, t_i) + (b - c)a_i - p_i, & \text{if the agent } i \text{ wins the object} \\ -ca_i - p_i, & \text{otherwise,} \end{cases}$$

where p_i denotes the total amount of payments made by the agent i within a mechanism (i.e. not including c).⁶ Note that p_i need not be equal to zero for loosing bidders. We assume that $V(s_i, t_i)$ is continuous and increasing in both arguments.

To explore the model, we need to extend the concepts of allocation, efficiency, and social surplus to this environment. An *outcome* is a vector consisting of the list of agents that invest and the identity of the agent that receives the object. *Social surplus* is the value of the object to the agent that gets the object minus the cost of investment taken by all agents: $S(\mathbf{a}) = V(s_j, t_j) + ba_j - \sum_{i=1}^N ca_i$, where j is the identity of the agent that receives the object, and $\mathbf{a} = (a_1, \dots, a_N)$ is a vector of agents' investment decisions. Outcomes can be ranked in terms of efficiency by comparing corresponding expected values of the social surplus. An outcome is *efficient* (first-best), if it yields the same expected social surplus as the maximum expected social surplus that can be achieved by the social planner, which observes all signals received by agents, orders agents to invest or not to invest in the second period, and, finally, assigns the object.

Example 2 *Our model includes, as a particular case, the situation where the second signal is a refinement of the first one. Indeed, let s_i, t_i be independent estimates of the object's value, taken*

⁶It is possible to extend our model to the case when the utility function takes the form $U_i(s_i, t_i, a_i)$, where a_i is continuous, and higher values of a_i makes a higher third-period signal t_i more likely. However, it would make exposition much more complex, while providing no new insight.

before and after the investment stage, and drawn from the same distribution. Then the agent's value function net of investment is $V(s_i, t_i) = \frac{1}{2}(s_i + t_i)$.

Example 3 This example illustrates that the auction with entry, where bidders get private information both before and after they make their entry decisions, is a particular case of our setup. Indeed, suppose that all signals s, t are distributed on a finite support, $W(s_i, t_i)$ is a function which is continuous and increasing in both arguments, and $D > 0$ is such that $\max_{s_i, t_i} W(s_i, t_i) < D$. Let $b = D$ and $V(s_i, t_i) = W(s_i, t_i) - D$. Agents that decided not to invest at the investment stage are not competing at the final stage, as their value of the object is below zero. Effectively, the cost of investment, c , becomes the entry cost.

3 Efficient Mechanism

We start with considering a benchmark case of the efficient mechanism for allocating the object that can be achieved by a social planner who knows all the private information available to bidders and controls their moves. Then we consider a mechanism that allocates the object efficiently in the incomplete information case.

3.1 Complete Information

Suppose that, upon observing the first-period signals, a social planner decides which agents should invest in the second period. Formally, there is a mapping of a vector of the first period signals \mathbf{s} into a vector of the second period actions $\mathbf{a} = \mathbf{a}(\mathbf{s})$, $\mathbf{a}_i = 1$ if agent i invests, and $\mathbf{a}_i = 0$ otherwise.⁷ At the end of the third period the social planner assigns the object, thus mapping a triplet of vectors $(\mathbf{s}, \mathbf{a}, \mathbf{t})$ into a number between 1 and N . The final assignment of the object is easily characterized. The social surplus maximization calls for assigning the object to the agent with the highest ex-post private value: if the efficient allocation assigns the object to the agent j , then for any $i \neq j$, we have $V(s_j, t_j) + ba_j \geq V(s_i, t_i) + ba_i$. Thus, assigning the object before agents have learned their final values of the object is likely to be inefficient: giving the object to the agent with the highest ex-post value is only necessary, but not sufficient for efficiency. The function $\mathbf{a}^*(\mathbf{s})$ describes the second period actions that maximize the expected social surplus. (Hereinafter mathematical expectation is taken with respect to the second signal $\mathbf{t} = (t_1, \dots, t_N)$, unless specified otherwise.)

$$\max_{\mathbf{a}} E[S(\mathbf{a})|\mathbf{s}] = \max_{\mathbf{a}} \left\{ E \left[\max_i \{V(s_i, t_i) + ba_i\} \right] - c \sum_{j=1}^N a_j \right\}.$$

It is useful to introduce a function $G_i(\mathbf{s}, \mathbf{a}_{-i})$ representing the difference in the expected social surplus that results from the agent i investing and not investing (keeping the investment decisions

⁷Though the social planner may assign mixed strategies to the agents, we show later that, almost surely, the social planner problem has a unique pure strategy solution. Consequently, we focus on pure strategies of the social planner.

of other agents unchanged):

$$G_i(\mathbf{s}, \mathbf{a}_{-i}) = E[S|\mathbf{s}, \mathbf{a}_{-i}, a_i = 1] - E[S|\mathbf{s}, \mathbf{a}_{-i}, a_i = 0]. \quad (1)$$

Since the social planner maximizes social surplus, the expected surplus in the above formula should be computed under assumption that after the third period, the social planner allocates the object to the agent with the highest value. The social planner faces the following trade-off: each additional agent's investment increases the expected private value of the agent who receives the object, but is associated with the cost of c .

In a world without an all-knowing and well-intentioned social planner, agents have to invest non-cooperatively. As a step towards an efficient mechanism for the main case, suppose that the first-period signals are common knowledge. We show that an efficient allocation can be achieved in a decentralized case, when bidders know each other's first-period signals. Part (ii) of Theorem 1 states that in this case, there exists an equilibrium outcome of a second price auction conducted at the end of the third period that yields an efficient allocation, the same allocation as the first best obtained by the social planner. The basic intuition is as follows: the expected increase in an agent's utility from taking an action is exactly equal to the change in the expected social surplus due to her action.⁸ A straightforward proof is relegated to the Appendix.

Theorem 1 (i) *For a given vector of the first-period private signals \mathbf{s} , there exists a threshold $r^* = r^*(\mathbf{s})$ such that the social planner assigns agents with the highest r^* first-period signals to invest.*⁹

(ii) *If first-period signals \mathbf{s} are public knowledge, there exists an efficient perfect Bayesian equilibrium of the game with a second price sealed bid auction conducted at the end of the third period. In this equilibrium, agents invest as if they were assigned by the social planner according to the strategy characterized in (i).*

Here and in the rest of the paper the term 'equilibrium' is reserved for a perfect Bayesian equilibrium. The efficient equilibrium described in Theorem 1 seems to be a natural focal point. However, the game has a coordination component: there are other perfect Bayesian equilibria that are not efficient. For example, if there are only two players, there might be multiple equilibria: e.g., one with the highest-ranked agent investing and the other abstaining, and another one with the second-ranked agent acting and the highest-ranked abstaining.

The above results cannot be naturally extended to the case when agents' signals are not symmetric or independent. The most simple example that highlights this point is as follows.

⁸The logic behind the result is similar to the one that insures efficient entry in McAfee and McMillan (1987) and Levin and Smith (1994). However, in these two papers, potential bidders are symmetric and have no private information prior to entry.

⁹ r^* is determined almost uniquely: The event that the expected social surplus is maximized by more than one action vector of the form $\mathbf{a}(r^*)$ and $\mathbf{a}(r^{**})$ where $r^* \neq r^{**}$ has zero probability.

Example 4 Suppose that realizations of first-period signals are $s_1 = 0$ and $s_2 = \varepsilon > 0$. Let signals t_1, t_2 be distributed identically, but not symmetrically (and not independently): $t_1, t_2 \in \{0, \mu, 2\}$, where $\mu < 1$. Joint probabilities are as follows: $\Pr(0, 0) = \Pr(0, 2) = \Pr(\mu, 0) = \Pr(\mu, \mu) = \Pr(2, \mu) = \Pr(2, 2) = \frac{1}{6}$ and $\Pr(0, \mu) = \Pr(\mu, 2) = \Pr(2, 0) = 0$. Then $\max\{t_1 + 1, t_2\} \succeq \max\{t_1, t_2 + 1\}$. Let b , the benefit from investing, be equal to 1. If $\varepsilon > 0$ is sufficiently small, $\max\{t_1 + 1, t_2 + \varepsilon\} \succeq \max\{t_1, t_2 + \varepsilon + 1\}$ as well, which means that maximization of the expected surplus requires the first agent (who has a lower first signal) to invest, and the second agent not to invest (for some $c > 0$).

To complete the example, suppose that signals s_1, s_2 are distributed independently and can take values $\{0, \varepsilon\}$ with equal probability. For the pairs $(0, \varepsilon)$ and $(\varepsilon, 0)$ of realized first-period signals, the distribution of the third-period signals is as described above (with the symmetric change for the case $(\varepsilon, 0)$). For pairs $(0, 0)$ and $(\varepsilon, \varepsilon)$, let signals $t_1, t_2 \in \{0, \mu, 2\}$ be distributed identically and independently.

3.2 Incomplete Information

Now we are ready to investigate the incomplete information case: the only difference with the previous section is that here bidders' signals regarding their private values (\mathbf{s} and \mathbf{t}) are observed privately. A mechanism consisting of an auction conducted after the third period no longer leads to an efficient allocation, since under such mechanism agents take second-period actions without knowledge of the private signals obtained by other players.¹⁰ (Obviously, an efficient allocation rule cannot always assign the final ownership of the object prior the end of the third period.)

We explicitly construct an efficient allocation mechanism, which consists of two rounds: The private information is revealed by signalling that takes place after the first-period-private-signals are observed; the ownership of the object is assigned in the second round that takes place after the third-period private signals are observed by bidders. Our mechanism does not require any commitment on the auctioneer's part. In contrast, the dynamic VCG mechanism by Bergemann and Valimaki (2000), though does ensure efficient allocation in this setup, is heavily reliant (as any VCG mechanism, see, e.g., Krishna, 2002) on the ability of the auctioneer to pay some bidders huge sums ex-post.

¹⁰For the sake of completeness, one can consider the no-signalling case, where an auction is conducted after the third period and no signaling takes place before the second period. (Note that cheap talk communication following the first stage is not credible because everybody has an incentive to exaggerate his signal.) To describe the symmetric equilibria of this game, one can show that there exists a unique constant s^* such that any agent acts if her first-stage value s_i is higher or equal to s^* , and abstains from acting otherwise. In the equilibrium, the expected number of actions is $N(1 - F_s(s^*))$. So, in some cases, there are too few actions, while in others there are too many. (cf. Example 1.) Also, there are a number of asymmetric equilibria, which cannot lead to an efficient allocation rule.

The Signalling Mechanism

1. At the end of the first period (after the private signals \mathbf{s} have been received by agents), all agents make simultaneous public announcements \hat{s}_i about their private values s_i . Also, each agent voluntarily selects a payment amount, $h_i = h_i(\hat{\mathbf{s}}) \geq 0$, that depends on the announcements of other agents, as well as on her own announcement, to make his announcement credible.¹¹
2. The second round takes place at the end of the third period, after agents observe their private signals t_i . The final ownership of the object is assigned using a second-price sealed-bid auction.

Theorem 2 below asserts that there exists a perfect Bayesian equilibrium of the Signalling Mechanism that yields an efficient outcome. The logic behind the result is as follows. First, if the first-period signals are revealed truthfully, the remaining subgame is identical to the game where first-period signals \mathbf{s} are common knowledge. Thus, in order to establish existence of an efficient allocation mechanism, it suffices to show that for some payment schedule $h_i(\hat{\mathbf{s}})$, truthful reporting is an equilibrium, when agents anticipate that the equilibrium characterized in Theorem 1 will be played in the remaining subgame. Agents are willing to pay in order to reveal their first period signals, because this information discourages other agents from taking actions, thus increasing the probability of winning for agent i and decreasing the expected price that she will pay for the object in the subsequent second price auction conditional on winning. To make truthful revelation possible, we need to demonstrate that the higher is the first period signal s_i received by agent i , the higher is that agent's relative willingness to pay in order to signal that her value of s_i is high.

Lemma 1, which is the main technical result in this Section, establishes an appropriate analog of the increasing-differences property for the payoffs in the subgame. Once this property is established, the existence of a payment schedule that is consistent with the incentive compatibility and individual rationality for all agents becomes a standard exercise.

Lemma 1 *Let $E[\pi_i(s_i, \hat{s}_i, s_{-i})]$ be bidder i 's expected pay-off gross of $h_i(\hat{\mathbf{s}})$, when her true private signal is s_i , while other agents believe that the vector of first-period private signals is (\hat{s}_i, s_{-i}) . For any s_{-i} and any $\hat{s}'_i > \hat{s}_i$, and any $s'_i > s_i$,*

$$E[\pi_i(s'_i, \hat{s}'_i, \mathbf{s}_{-i})] - E[\pi_i(s'_i, \hat{s}_i, \mathbf{s}_{-i})] \geq E[\pi_i(s_i, \hat{s}'_i, \mathbf{s}_{-i})] - E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})]. \quad (2)$$

Proof of Lemma 1. First, we claim that for any i , and for any \mathbf{s} and $\tilde{\mathbf{s}}$ such that $\mathbf{s}_{-i} \leq \tilde{\mathbf{s}}_{-i}$ and $s_i = \tilde{s}_i$, $a_i^*(\tilde{\mathbf{s}}) \leq a_i^*(\mathbf{s})$. Indeed, let $X_i(\mathbf{s}) = V(s_i, t_i)$ and $Z_i(\mathbf{s}) = \max_{j \neq i} \{V(s_j, t_j) + ba_j^*\}$. For any number (function) x , denote $x^+ = \max\{x, 0\}$.

Suppose that $a_j^*(s_i, \mathbf{s}_{-i}) \leq a_j^*(s_i, \tilde{\mathbf{s}}_{-i})$ for all $j \neq i$. Then

$$X_i(s_i, \mathbf{s}_{-i}) - Z_i(s_i, \mathbf{s}_{-i}) \succeq X_i(s_i, \tilde{\mathbf{s}}_{-i}) - Z_i(s_i, \tilde{\mathbf{s}}_{-i})$$

¹¹These payments can be either money burning or transfers to a third party, or even to the seller.

by Lemma A (1). To prove that $g_i(\mathbf{s}) \geq g_i(\tilde{\mathbf{s}})$, recall that

$$g_i(\mathbf{s}) = E[X_i(\mathbf{s}) + b - Z_i(\mathbf{s})]^+ - E[X_i(\mathbf{s}) - Z_i(\mathbf{s})]^+,$$

and then apply Lemma A (3) to prove the claim. By definition, $g_i(\mathbf{s}) \geq g_i(\tilde{\mathbf{s}})$ implies that $a_i^*(\tilde{\mathbf{s}}) \leq a_i^*(\mathbf{s})$.

It is enough to consider the case of $a_j^*(s_i, \mathbf{s}_{-i}) \leq a_j^*(s_i, \tilde{\mathbf{s}}_{-i})$ for all $j \neq i$. Indeed, if a switch from 1 to 0 occurred with an agent that ends up higher than i as a result of an increase from \mathbf{s}_{-i} to $\tilde{\mathbf{s}}_{-i}$, then it is definite that $a_i^*(\tilde{\mathbf{s}}) = 0$, and thus $a_i^*(\tilde{\mathbf{s}}) \leq a_i^*(\mathbf{s})$ for any $a_i^*(\mathbf{s})$. Otherwise (if a change have occurred with an agent ranked lower than the agent i), $a_i^*(\mathbf{s}) = 1$.

Second, we claim that the function g_i increases with s_i . The first claim shows, in particular, that if s_i increases, while \mathbf{s}_{-i} is constant, the number of agents acting (weakly) decreases. Thus, the random variable $X_i(s_i, \mathbf{s}_{-i}) - Z_i(s_i, \mathbf{s}_{-i})$ raises in terms of stochastic dominance, and Lemma A (3) applies.

Now we shall prove that

$$E[\pi_i(s'_i, \hat{s}'_i)] - E[\pi_i(s_i, \hat{s}'_i)] \geq E[\pi_i(s'_i, \hat{s}_i)] - E[\pi_i(s_i, \hat{s}_i)],$$

which is equivalent to (2).

Define $X_i = V(s_i, t_i)$, $X'_i = V(s'_i, t_i)$, $Y_i = \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\hat{s}_i, \mathbf{s}_{-i})\}$, and $Y'_i = \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\hat{s}'_i, \mathbf{s}_{-i})\}$.

$$\begin{aligned} E[\pi_i(s'_i, \hat{s}_i)] - E[\pi_i(s_i, \hat{s}_i)] &= E[X'_i - Y_i]^+ - E[X_i - Y_i]^+, \\ E[\pi_i(s'_i, \hat{s}'_i)] - E[\pi_i(s_i, \hat{s}'_i)] &= E[X'_i - Y'_i]^+ - E[X_i - Y'_i]^+, \end{aligned}$$

and so it remains to prove that

$$E[X'_i - Y'_i]^+ - E[X_i - Y'_i]^+ \geq E[X'_i - Y_i]^+ - E[X_i - Y_i]^+.$$

The two claims proved above yield that $X'_i \succeq X_i$ and $Y_i \succeq Y'_i$. Using Lemma A3 (for each non-negative constant) completes the proof. ■

In the above Lemma, $E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})]$ is the expected pay-off of agent i in the mechanism described in Section 3.1, when the first period private signals are given by (s_i, \mathbf{s}_{-i}) and player i plays the best response to the action profile of players $-i$ given by $\mathbf{a}(r^*(\hat{s}_i, \mathbf{s}_{-i}))$. (The action profile $\mathbf{a}(r^*(\hat{s}_i, \mathbf{s}_{-i}))$ is characterized in Theorem 1.) Essentially, $E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})]$ is the pay-off received by agent i in the subgame computed under an assumption that all first round announcements are believed to be truthful, and that agent i reported \hat{s}_i , while her true private value is s_i . Lemma 1 states that the same change in announcement (from \hat{s}_i to \hat{s}'_i) brings more in expected surplus to the agent with relatively high true signal, s'_i . Note that $E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})]$ is not the same as the expected utility of agent i , because it does not include the payments h_i made in the first round of the mechanism. The agent's utility is given by $E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})] - h_i$. Thus, truthfully reporting s_i

is consistent with an equilibrium, if there exists a payment schedule $h_i(\hat{s}_i, \mathbf{s}_{-i})$ such that incentive compatibility and individual rationality constraints are satisfied. Namely, for any agent i and all $(s_i, \hat{s}_i, \mathbf{s}_{-i})$ the payments should satisfy the following conditions:

$$E[\pi_i(s_i, s_i, \mathbf{s}_{-i})] - h_i(s_i, \mathbf{s}_{-i}) \geq E[\pi_i(s_i, \hat{s}_i, \mathbf{s}_{-i})] - h_i(\hat{s}_i, \mathbf{s}_{-i}) \quad (\text{IC})$$

$$E[\pi_i(s_i, s_i, \mathbf{s}_{-i})] - h_i(s_i, \mathbf{s}_{-i}) \geq E[\pi_i(s_i, \hat{s}_i = 0, \mathbf{s}_{-i})] \quad (\text{IR})$$

Finding some $h_i(\hat{s}_i, \mathbf{s}_{-i})$ that satisfies the above constraints is sufficient for proving the claim of Theorem 2. Still, before proceeding to Theorem 2, we need to introduce one more definition. Consider the efficient allocation rule characterized in Theorem 1. It implies that for any vector of the first period private signals \mathbf{s}_{-i} , there exists a sequence $0 = \bar{s}_i(k_i^*) \leq \bar{s}_i(k_i^* - 1) \leq \dots \leq \bar{s}_i(1) \leq \bar{s}_i(0) < \infty$, where $\bar{s}_i(k)$ is defined to be the minimal type of i such that exactly k highest-ranked agents (different from the agent i herself) invest in the subgame equilibrium described in Theorem 1. Let $k_i^* = k_i^*(0, \mathbf{s}_{-i})$ be the number of agents investing, when i has the lowest possible type (zero). Within each segment described above, an agent's i report is irrelevant to the other agents' decisions on whether or not to invest. Let $\mathbf{a}(m)$ denote the vector of actions, where the agents with the highest m first-period signals invest, while the other $N - m$ agents not invest. Note that $\mathbf{a}(m)$ is a function of the vector of first-period signals \mathbf{s} .

Theorem 2 *The following payment schedules are consistent with the incentive compatibility and individual rationality conditions (IC and IR, respectively). For any i ,*

$$h_i(\hat{s}_i, \hat{\mathbf{s}}_{-i}) = 0, \text{ whenever } \bar{s}_i(k_i^*) \leq \hat{s}_i \leq \bar{s}_i(k_i^* - 1), \quad (3)$$

$$h_i(\hat{s}_i, \hat{\mathbf{s}}_{-i}) = h_i(\bar{s}_i(k), \hat{\mathbf{s}}_{-i}) + E[\pi_i(\bar{s}_i(k), \mathbf{a}(k))] - E[\pi_i(\bar{s}_i(k), \mathbf{a}(k+1))],$$

whenever $\bar{s}_i(k) < \hat{s}_i \leq \bar{s}_i(k-1)$, $k < k_i^$. Thus, there exists an efficient perfect Bayesian equilibrium in the Signalling Mechanism.*

Proof. Let s_i be the true agent's i first-period signal, and consider k such that $\bar{s}_i(k) < \hat{s}_i \leq \bar{s}_i(k-1)$. Since $\hat{\mathbf{s}}_{-i}$ is fixed throughout the argument, we suppress the notation. Truthful reporting brings the expected utility of

$$E[\pi_i(s_i, \mathbf{a}(k))] - h_i(s_i) = E[\pi_i(s_i, \mathbf{a}(k))] - h_i(\bar{s}_i(k)) - E[\pi_i(\bar{s}_i(k), \mathbf{a}(k))] + E[\pi_i(\bar{s}_i(k), \mathbf{a}(k+1))].$$

First, we prove that the agent i has no incentives to under-report her first-period signal, i.e. to report $\hat{s}_i < s_i$. Consider incentives the agent i with the first-period signal $\bar{s}_i(k)$ faces. For any ε such that $\bar{s}_i(k) - \bar{s}_i(k+1) > \varepsilon > 0$, she is indifferent between reporting $\bar{s}_i(k)$ and reporting $\bar{s}_i(k) - \varepsilon$. Indeed, the payment is the same and the number of acting rivals is the same ($k+1$). The condition (2) assures that if the agent with $\bar{s}_i(k)$ is indifferent between reporting $\bar{s}_i(k)$ to reporting $\bar{s}_i(k) - \varepsilon$, then the agent with $s_i > \bar{s}_i(k)$ (weakly) prefers reporting $\bar{s}_i(k)$ to reporting $\bar{s}_i(k) - \varepsilon$. Thus, \hat{s}_i

can not be less than $\bar{s}_i(k)$. (To rule out reports below $\bar{s}_i(k+1)$, one can consider incentives the $\bar{s}_i(k+1)$ -agent faces.) It remains to show that \hat{s}_i (weakly) exceeds $\bar{s}_i(k)$. So, we need to prove that $E[\pi_i(s_i, \mathbf{a}(k))] - h_i(\bar{s}_i(k)) - E[\pi_i(\bar{s}_i(k), \mathbf{a}(k))] + E[\pi_i(\bar{s}_i(k), \mathbf{a}(k+1))] \geq E[\pi_i(s_i, \mathbf{a}(k+1))] - h_i(\bar{s}_i(k))$, or equivalently,

$$E[\pi_i(s_i, \mathbf{a}(k))] - E[\pi_i(s_i, \mathbf{a}(k+1))] \geq E[\pi_i(\bar{s}_i(k), \mathbf{a}(k))] - E[\pi_i(\bar{s}_i(k), \mathbf{a}(k+1))],$$

but this is true by (2). Since the agent i having the signal s_i is indifferent between reporting s_i and reporting any signal that is larger than $\bar{s}_i(k)$ and does not exceed s_i , the proof that the agent i has no incentives to under-report her signal is complete.

The proof that there is no incentives to over-report the signal is analogous. The s_i -agent is indifferent between reporting the true signal and reporting $\bar{s}_i(k-1)$. Indeed, the mechanism assumes that the agents with reports s_i and $\bar{s}_i(k-1)$ pay the same amount. Now, for any ε such that $\bar{s}_i(k-2) - \bar{s}_i(k-1) > \varepsilon > 0$, the agent with $\bar{s}_i(k-1)$ is indifferent between reporting the true signal and reporting $\bar{s}_i(k-1) + \varepsilon$. To see this, note that

$$\begin{aligned} E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k))] - h_i(\bar{s}_i(k-1)) &= E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k-1))] - h_i(\bar{s}_i(k-1)) \\ &\quad - E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k-1))] + E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k))]. \end{aligned}$$

By (2),

$$E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k-1))] - E[\pi_i(\bar{s}_i(k-1), \mathbf{a}(k))] \geq E[\pi_i(s_i, \mathbf{a}(k-1))] - E[\pi_i(s_i, \mathbf{a}(k))].$$

Thus, if the $\bar{s}_i(k-1)$ is indifferent between reporting the truth and reporting $\bar{s}_i(k-1) + \varepsilon$, the s_i -agent (weakly) prefers to report $\bar{s}_i(k-1)$ (which is pay-off equivalent to reporting the truth), than to report $\bar{s}_i(k-1) + \varepsilon$. To show, that \hat{s}_i would not exceed $\bar{s}_i(k-2)$, one should consider the incentives the $\bar{s}_i(k-2)$ -agent faces, etc. Therefore, the agent i has no incentives to over-report her first-period signal.

Now, if agents in the set $-i$ report their type truthfully, $\hat{\mathbf{s}}_{-i} = \mathbf{s}_{-i}$, the payment scheme for the agent i given by (3) induces her to report her type truthfully, $\hat{s}_i = s_i$. Lemma 1 ensures that the above payment schedule induces truthful reporting by agent i , provided that all other agents' reports are truthful. The beliefs supporting the equilibrium in the signalling stage are straightforward: if a payment by an agent i is defined by (3), then the agents first-period signal is perceived to lie within the respective range. In the subgame that starts after the first-period signals are revealed, agents play according to the strategies described in Theorem 1. ■

3.3 Ex-ante Efficient Equilibrium

The mechanism described above provides an ex-post efficient ex-post equilibrium (Cremer and McLean, 1985), where no agent regrets her announcement after learning the other agents' types.

In such an equilibrium, agents' payments may depend on the other agents' announcements. Below we show that the Signalling Mechanism described in the previous section also has an ex-ante efficient separating equilibrium. In this equilibrium, agents make no announcements (or make uninformative announcements) in the cheap talk stage of the mechanism. They simultaneously make publicly observable payments H_i ; an agent decides on the payment size without knowing the private signals of other agents. We show that there exists a fully separating equilibrium where there is a unique payment corresponding to each private signal s_i . Consequently, agents no longer need to make announcements, because the announcements of their private signals are revealed in the size of payments they make.

Theorem 3 *There exists an efficient equilibrium in the Signalling Mechanism, where agents simultaneously make payments $H_i(s_i)$ that depend only on their private information s_i . Equilibrium payments are given by $H_i(\hat{s}_i) = E_{\mathbf{s}_{-i}}[h_i(\hat{s}_i, \mathbf{s}_{-i})]$, where $h_i(\hat{s}_i, \mathbf{s}_{-i})$ are equilibrium payments defined in Theorem 2.*

Proof of Theorem 3. According to Theorem 1, an efficient allocation can be obtained, when the first period signals s_i are common knowledge. It remains to show that the signalling mechanism proposed above is incentive compatible when an efficient equilibrium is chosen in the subgame following the signalling stage. It suffices to observe that

$$s_i \in \arg \max_{\hat{s}_i} \{E[\pi_i(\hat{s}_i, \mathbf{s}_{-i})] - h_i(\hat{s}_i, \mathbf{s}_{-i})\}$$

for any \mathbf{s}_{-i} and any s_i , and take sum over all \mathbf{s}_{-i} .

Then note that all $h_i(\hat{s}_i, \mathbf{s}_{-i})$ and thus the function $H_i(s_i)$ increases in the bidder's i first-period signal s_i . This allows to use $H_i(s_i)$ to report the true value of s_i . Beliefs are straightforward. ■

Although existence of an ex-ante equilibrium follows directly from existence of an ex-post equilibrium, it is a useful result. It shades some light on the maneuvering that bidders often make prior to an auction: for example, firms preparing to participate in a large-scale privatization auction or competing for a procurement contract might engage in costly signalling in order to discourage potential rivals. For example, consider the story of selling the Los Angeles license in 1995 broadband auction for mobile-phone licenses. One bidder, Pacific Telephone, possibly started with a higher private value than other bidders due to experience in California market and possible synergies between its wireline and wireless businesses. There was a number of important decisions (actions) that each bidder had to make before the auction for Los Angeles license, these included forming alliances, making investments and formulating strategies for other markets. It appears that Pacific Telephone signaled to other bidders (and would-be bidders) that it anticipates winning California. Pacific Telephone made public statements like 'If somebody takes California away from us, they'll never make any money'.¹² To make these statement credible, Pacific Telephone made investments

¹² *Wall Street Journal*, October 31, 1994.

that were of little value without winning Los Angeles license¹³. As a result, some potential bidders (including the industry giants such as Bell Atlantic, GTE, and MCI) were discouraged from participating in the auction. (Thus failing to undertake an action, in our interpretation). In fact, GTE and Bell Atlantic took actions that made them ineligible for the auction. As a result, revenues were quite low compared to initial estimates.¹⁴ Applying the logic of our model highlights the importance of signalling that discourages competitors from taking actions that increase the value of the prize for them.

4 A Robust Auction

The efficient mechanism described in the previous section can be viewed as a two-stage auction. The reporting stage, where agents simultaneously make payments that reveal their types, might be replaced with a sealed-bid all-pay auction, where the object being sold is worth nothing (zero). Theorem 3 established existence of an efficient perfect Bayesian equilibrium of this two-stage auction. Unfortunately, this is not a unique equilibrium: a pooling equilibrium, where everybody bids zero in the signalling stage, is a natural focal point. Nevertheless, introducing a possibility of a small inefficiency into the auction design can force bidders to coordinate on an efficient separating equilibrium. Also, this new mechanism allows seller to capture the signalling costs of bidders. The sacrifice that must be made in order to gain robustness and capture signalling costs is an arbitrarily small loss in efficiency. We will refer to such a mechanism as an ε -coupon *mechanism*. We start with describing an ε -coupon mechanism and then proceed to establish efficiency properties of this mechanism in Theorem 4.

The ε -Coupon Mechanism

1. The first (reporting) round takes place at the end of the first period (after the private signals \mathbf{s} have been received by agents, but before agents take actions). In this round one, coupon is sold via

¹³Some of the investments made by Pacific Telephone might be interpreted as actions and others as signals. Essentially, running a PR campaign aimed at signaling that Pacific Telephone is determined to win Los Angeles license can be interpreted as signaling. In contrast, making unobservable arrangements made to expedite creation of the wireless service in Southern California can be interpreted as an action.

¹⁴This is not the only possible interpretation of the 1995 auction for Los Angeles licence. Klemperer (2000) considers the history of this auction and suggests that the winner's curse played an important role because the winner's curse is particularly powerful in auctions where one bidder has an advantage. For a theoretical argument that uses this logic, see also Bulow, Huang, and Klemperer (1999). The outcome of that auction was probably determined by a constellation of a large number of factors. Revenue in the auction for Los Angeles licence were low in comparison with spectrum auction in Chicago; however, it is not clear if asymmetry among bidders and the winner curse were more severe in California.

all pay sealed bid auction.¹⁵ All bids are announced at the end of the round. The coupon sold in the signalling round entitles its owner to a discount of size ε for the price in the final auction. The discount coupon is not-transferable: only the winner of the final auction can benefit from having the coupon.

2. The second round (final) auction takes place at the end of the third period, after agents observe private signals \mathbf{t} . In the second round the ownership of the object is assigned using a second price sealed bid auction. If the highest bidder in the final round is the owner of the ε -coupon, then she pays the second highest bid minus ε .

Formally, there are two rounds and three decision nodes in an ε -coupon mechanism. At the first decision node, agents make bids in an all-pay auction, i.e. the i 's actions space is $\{H_i | H_i \geq 0\}$. The information set of agent i at the first decision node is given by s_i . The first round strategy is described by the probability distribution $\rho_i(\cdot; s_i)$ over the set of pure strategies $\{H_i | H_i \geq 0\}$. At the second decision node, agents make a decision to invest or not to invest. The information set of agent i at the second decision node is given by $(s_i, H_i, \mathbf{H}_{-i}, \mathbf{w})$, where \mathbf{w} is an N -dimensional vector with $w_k = 1$ if the agent k won the coupon in the all pay auction, and $w_k = 0$ otherwise. (There is a unique vector \mathbf{w} consistent with vector of payments \mathbf{H} , unless there is a tie). The probability that agent i invests ($a_i = 1$) is denoted by $\lambda_i = \lambda_i(s_i, H_i, \mathbf{H}_{-i}, \mathbf{w})$. At the third decision node, agents submit bids in the second price sealed-bid auction. At this node, the information sets are $(s_i, H_i, \mathbf{H}_{-i}, \mathbf{w}, a_i, t_i)$. It is well known that in an equilibrium in weakly dominant strategies of a private value Vickrey auction bidders bid their true values. Thus, equilibrium bids are given by $V(s_i, t_i) + ba_i + \varepsilon w_i$.

Clearly, an ε -coupon mechanism has multiple equilibria. Some of these equilibria are highly implausible. In order to rule out such equilibria we introduce a restriction on strategies in the spirit of ‘intuitive’ criteria such as D1 of Cho and Kreps (1987) or stability of Kohlberg and Mertens (1986). A strategy of agent j is *monotonic*, if two vectors \mathbf{H}_{-j} and \mathbf{H}'_{-j} differ only in a component i so that if $H_i > H'_i$, then $p_j(s_j, H_j, \mathbf{H}_{-j}, \mathbf{w}) \leq p_j(s_j, H_j, \mathbf{H}'_{-j}, \mathbf{w}')$. In other words, a monotonic strategy of agent j assumes that for any history, the probability that agent j invests is non-increasing in the size of the payment that some other agent j , $i \neq j$ makes in the signalling stage.

The requirement that the equilibrium strategies are monotonic rules out the ‘bizarre’ equilibrium, where all agents bid zero in the signalling stage and an agent who bids a positive amount is perceived to be of the lowest type. Basically, there are two reasons why an equilibrium strategy may not be monotonic: First, perverse beliefs may sustain an equilibrium in strategies that are

¹⁵In an all pay sealed bid auction every agent submits a sealed bid. All agents have to pay the amount of their bids regardless of whether or not they won the object. The agent with the highest bid receives the object. (In case of a tie the winner is randomly chosen from the set of highest bidders.) Fullerton and McAfee (1999) use an all-pay auction in their ‘contestant selection auction’.

not monotonic. An example of such ‘unnatural’ beliefs is as follows: The more an agent bids for a discount coupon, the lower is her perceived s_i . Obviously, this is counter-intuitive: the higher is an agent’s s_i , the more she values the discount coupon. The second possibility stems from coordination aspect of the game. If bids in the signalling stage are used as coordination devices for selecting a perfect Bayesian equilibrium in the remaining subgame, an equilibrium resulting from these beliefs may include strategies that are not monotonic.

Theorem 4 *In an ε -coupon mechanism, there exists a unique symmetric perfect Bayesian equilibrium in monotonic strategies. The probability that the equilibrium yields an efficient outcome converges to one as $\varepsilon \rightarrow 0$.*

Before giving a formal proof to the results, let us sketch the intuition why a pooling equilibrium where everybody bids zero for the coupon is not a perfect Bayesian equilibrium in monotonic strategies. Indeed, if everybody bids zero, the coupon can be purchased for an arbitrarily small amount. Thus, the pooling equilibrium is sustainable only if bidders are discouraged from bidding a positive amount by their beliefs that a positive bid would encourage other bidders to invest more aggressively in the action stage. However, this belief is inconsistent with strategies being monotonic. The same argument applies to any partially pooling equilibrium. We show that there are no equilibria in mixed strategies, because the willingness to pay for the discount is an increasing function of the bidder’s signal. The efficiency result for an ε -coupon equilibrium follows from Theorem 3 that establishes that for $\varepsilon = 0$, there exists an efficient symmetric equilibrium. To prove asymptotic efficiency of a robust equilibrium, we show that when ε approaches 0, the robust equilibrium converges to the equilibrium described in Theorem 3.¹⁶

Proof of Theorem 4. The proof of existence follows the pattern of the proof of Theorem 2. Construction of an ex-ante equilibrium in the previous section used existence of an ex-post equilibrium in a mechanism, where payments that make announcements credible are allowed to be functions of announcements. Here, we use the same idea. As an intermediate step, consider a mechanism, where the discount is not auctioned off using an all-pay auction. Instead, bidders announce their types in the reporting stage (much like in the mechanism described in Section 3). After the announcement, bidders make payments $h_i(\widehat{s}_i, \widehat{\mathbf{s}}_{-i})$ to make the announcement credible, and the bidder with the highest announced s_i receives the ε -discount.

1. After each agent privately learns s_i , all agents simultaneously announce their types in the cheap talk stage. Afterwards, each agent must make a payment of $h_i(\widehat{s}_i, \widehat{\mathbf{s}}_{-i})$. The agent with the highest first period announcement \widehat{s} receives the discount coupon (ties are broken using a lottery). Agents take action after observing announcements $\widehat{\mathbf{s}}$.

2. After the third period signals \mathbf{t} are revealed, the object is sold via a second-price sealed-bid auction.

¹⁶ Also, if the ε -efficient mechanism yields an inefficient outcome, efficiency losses are of magnitude ε .

We shall show that there exists a payment schedule $h_i(\widehat{s}_i, \widehat{\mathbf{s}}_{-i})$ such that truthful reporting supported by paying $h_i(\widehat{s}_i, \widehat{\mathbf{s}}_{-i})$ is an ex-post equilibrium. The private value of the bidder with the highest first-period signal is essentially boosted by the amount equal to the discount ε . Let $\widetilde{\mathbf{s}}(\widehat{\mathbf{s}})$ be a vector of ‘adjusted’ private value signals, where $\widetilde{s}_i = s_i + \varepsilon$ if $\widehat{s}_i > \widehat{s}_j = s_j$ for all $j \neq i$, and $\widetilde{s}_i = s_i$ otherwise. Assuming $\widehat{\mathbf{s}}_{-i} = \mathbf{s}_{-i}$, we study the i th agent incentives to misreport the true signal s_i . If all equilibrium reports \widehat{s}_i are truthful, then the subgame after ε discount is assigned is identical to the game considered in Section 3.1. The equilibrium expected pay-off of agent i of the subgame, which does not include h_i , is denoted by $E[\widetilde{\pi}_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})]$. One can express $E[\widetilde{\pi}_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})]$ in terms of $E[\pi_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})]$ (defined in Lemma 1) using ‘adjusted’ private signals. Let $\widehat{\widetilde{\mathbf{s}}}$ denote a vector of perceived ‘adjusted’ signals of agents; the i th component of $\widehat{\widetilde{\mathbf{s}}}$ is $\widehat{\widetilde{s}}_i = \widehat{s}_i(s_i, \widehat{s}_i, \widehat{\mathbf{s}}_{-i}) = \widetilde{s}_i + (s_i - \widehat{s}_i)$. That is, $\widetilde{\mathbf{s}}$ is a vector of ‘adjusted’ private value signals and $\widehat{\widetilde{\mathbf{s}}}$ is public perception about $\widetilde{\mathbf{s}}$. Now we can write $E[\widetilde{\pi}_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})] = E[\pi_i(\widetilde{\mathbf{s}}_i, \widehat{\widetilde{\mathbf{s}}}_i, \widetilde{\mathbf{s}}_{-i})]$.

To prove that a separating equilibrium exists, we need to formulate an increasing-differences condition similar to equation (2). We formally state this claim as follows:

For any $N - 1$ -tuple of truthful reports \mathbf{s}_{-i} , and any $s'_i \geq \widehat{s}_i$, $s'_i \geq s_i$,

$$E[\widetilde{\pi}_i(s'_i, \widehat{s}'_i, \mathbf{s}_{-i})] - E[\widetilde{\pi}_i(s'_i, \widehat{s}_i, \mathbf{s}_{-i})] \geq E[\widetilde{\pi}_i(s_i, \widehat{s}'_i, \mathbf{s}_{-i})] - E[\widetilde{\pi}_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})]. \quad (4)$$

To prove the claim, we need to consider three cases: (a) the agent wins the ε discount if she makes announcement \widehat{s}'_i but not \widehat{s}_i ; (b) an agent wins the discount for either announcement \widehat{s}'_i or \widehat{s}_i ; (c) neither \widehat{s}'_i , nor \widehat{s}_i are high enough to win the discount.

For (b) and (c), equation (4) follows immediately from Lemma 1. It remains to show that it also holds for the case (a). Denote $\mathbf{s}_{-i} = (s_{-i}^m, \mathbf{s}_{-i}^{-m})$, where s_{-i}^m is the largest component of the vector \mathbf{s}_{-i} and \mathbf{s}_{-i}^{-m} is an $N - 2$ -dimensional vector that consists of all components of vector \mathbf{s}_{-i} other than its largest component s_{-i}^m . Applying the new notation, one gets $E[\widetilde{\pi}_i(s_i, \widehat{s}_i, \mathbf{s}_{-i})] = E[\pi_i(s_i, \widehat{s}_i, s_{-i}^m, \mathbf{s}_{-i}^{-m})]$. In case (a), we have $\widetilde{\mathbf{s}}_{-i}^{-m} = \mathbf{s}_{-i}^{-m}$. Therefore, one can re-write (4) as follows:

$$E[\pi_i(s'_i + \varepsilon, \widehat{s}'_i, s_{-i}^m)] - E[\pi_i(s'_i, \widehat{s}_i, s_{-i}^m + \varepsilon)] \geq E[\pi_i(s_i + \varepsilon, \widehat{s}'_i, s_{-i}^m)] - E[\pi_i(s_i, \widehat{s}_i, s_{-i}^m + \varepsilon)]. \quad (5)$$

Let

$$\begin{aligned} X &= V(s_i, t_i) + \varepsilon - \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\widehat{s}'_i, s_{-i}^m, \mathbf{s}_{-i}^{-m})\}, \\ X' &= V(s_i, t_i) + \varepsilon - \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\widehat{s}_i, s_{-i}^m + \varepsilon, \mathbf{s}_{-i}^{-m})\}, \\ Y &= V(s'_i, t_i) + \varepsilon - \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\widehat{s}'_i, s_{-i}^m, \mathbf{s}_{-i}^{-m})\}, \\ Y' &= V(s'_i, t_i) + \varepsilon - \max_{j \neq i} \{V(s_j, t_j) + ba_j^*(\widehat{s}_i, s_{-i}^m + \varepsilon, \mathbf{s}_{-i}^{-m})\} \end{aligned}$$

We know that $X' \succeq X$, $Y' \succeq Y$. Then

$$\begin{aligned} E[\pi_i(s'_i + \varepsilon, \widehat{s}'_i, s_{-i}^m)] - E[\pi_i(s_i + \varepsilon, \widehat{s}'_i, s_{-i}^m)] &= E[Y]^+ - E[X]^+, \\ E[\pi_i(s'_i, \widehat{s}_i, s_{-i}^m + \varepsilon)] - E[\pi_i(s_i, \widehat{s}_i, s_{-i}^m + \varepsilon)] &= E[Y']^+ - E[X']^+. \end{aligned}$$

Using Lemma A3 (from the Appendix) completes the proof of (5).

Since (5) holds, there exists an ex-post separating equilibrium in the “intermediate mechanism”. Using existence of an ex-post equilibrium, we can apply the same argument as in the proof of Theorem 3 to establish existence of ex-ante separating signalling mechanism, where agents make signalling payments that are strictly increasing in their signals.

Now we shall prove that any perfect Bayesian equilibrium in monotonic strategies is unique, fully separating, and ‘almost efficient’. In an equilibrium, the probability of any particular bid with value H in the signalling stage is zero. Indeed, if there is a positive mass of agents that plays some H_{mass} with positive probability, then there is a positive probability of a tie. Then an agent playing H_{mass} can increase the likelihood of winning the discount $\varepsilon > 0$ by increasing her bid by an infinitesimal amount. Since the strategies are monotonic, none of the agents would increase their likelihood of taking actions. Thus, such a deviation would be profitable.

Probability that players in the set $-i$ take actions is denoted here as λ_{-i} . Let $\Pi(s_i, \lambda_{-i}, \mathbf{s}_{-i})$ denote the pay-off of player i in the subgame after signalling payments H ’s are sunk. We want to show that if $\lambda_{-i} \geq \lambda'_{-i}$ then for every $s'_i > s_i$ we have

$$\Pi(s'_i, \lambda_{-i}, \mathbf{s}_{-i}) - \Pi(s_i, \lambda'_{-i}, \mathbf{s}_{-i}) \leq \Pi(s'_i, \lambda_{-i}, \mathbf{s}_{-i}) - \Pi(s_i, \mathbf{p}'_{-i}, \mathbf{s}_{-i}). \quad (6)$$

Thus, any decrease in “final” private values of player in the set $-i$ is more valuable for player i with a larger first period private signal. Inequality (6) follows from the proof of Lemma 1.

Let us show that all equilibria in monotonic are separating. In such an equilibrium, actions taken by players depend on their private signals and the announcements of other players; one can write $\lambda_{-i} = \lambda_{-i}(\mathbf{s}_{-i}, \mathbf{H}_{-i}, H_i)$ and $\lambda'_{-i} = \lambda_{-i}(\mathbf{s}_{-i}, \mathbf{H}_{-i}, H'_i)$. (From above, we can conclude that a tie is a measure zero event; and so have no impact on expected payoffs.) For monotonic strategies $\lambda_{-i} \geq \lambda'_{-i}$ for $H'_i > H_i$ (the inequality holds for all components). Inequality (6) implies that $H(s)$ is weakly increasing in s . Now we can conclude that $H(s)$ is strictly increasing in s almost everywhere.

Let us show that in an equilibrium, $\lambda_i(\mathbf{H}_{-i}, H_i(s_i), s_i)$ is non decreasing in s_i . Indeed, $\lambda_{-i} = \lambda_{-i}(\mathbf{s}_{-i}, \mathbf{H}_{-i}, H_i)$ is weakly decreasing in H_i . According to single crossing condition, if an agent with a first period signal s_i invests with positive probability $\lambda_i(\mathbf{H}_{-i}, H_i(s_i), s_i) > 0$, then any agent with a signal $s'_i > s_i$ strictly prefers to invest, and $\lambda_i(\mathbf{H}_{-i}, H_i(s'_i), s'_i) = 1$. Therefore, there exists a unique equilibrium in the subgame which is consistent with a monotonic strategy profile. In this equilibrium, all agents with private values exceeding some critical value $s^*(\mathbf{H})$ invest. From the previous paragraph and Theorem 1, it follows that ε -coupon mechanism yields an efficient allocation with probability converging to one as ε converges to zero.

To establish uniqueness of the robust equilibrium, we use a standard argument (e.g., Klemperer, 1999). Condition (6) implies that $\frac{dH(s)}{ds}$ is the same in any robust equilibrium. Above, we showed that there is a unique robust equilibrium in the subgame following the all-pay auction. It remains

to show that $H(0) = 0$. Suppose otherwise, say $H(0) = H_0 > 0$. For a player with $s_i = 0$, $H(0) = 0$ is a profitable deviation: Indeed, after this she does not change the perception of her type (she is correctly perceived to have $s_i = 0$). It was demonstrated that in a robust equilibrium each player either invests with probability one or zero (except perhaps for a set of measure zero). Thus, the deviation can only cause other players to increase the probability with which they invest; however, given the set of players that invest, none of the players that do not invest in a robust equilibrium would choose to invest. ■

Let us now consider an example illustrating that the all pay auction part of the ε -coupon mechanism is crucial for ensuring that any perfect Bayesian equilibrium in monotonic strategies is separating and nearly efficient.

Example 5 *Suppose the all-pay auction is replaced with a second-price sealed-bid auction. When a sufficiently small discount is auctioned off via a second price auction, the following inefficient pooling equilibrium is robust: all agents bid ε for the discount of size ε . Indeed, we need to specify beliefs that support this equilibrium. If an agent deviates by bidding less than ε , she is perceived to have the lowest possible signal s_i . Thus, there are no incentives to bid less than ε , provided that ε is sufficiently small. If an agent bids more than ε , the beliefs of other agents about her type are the same as if she bids ε . Thus, bidding more than ε is a bad strategy: If there are N agents bidding ε each in a second-price auction, each of them has a $\frac{1}{N}$ chance of getting the discount. The winner of the discount “envies” the bidders who did not win the discount, and thus do not have to pay anything in the signalling stage. By bidding more than ε , an agent insures that she wins the discount and will have to pay for it, thus, making herself worse off. In contrast, there are no robust pooling equilibrium of the ε -coupon mechanism (by Theorem 4). For instance, if all agents bid ε for the discount, bidding slightly more than ε is a profitable deviation.*

5 Conclusion

We put forth an efficient mechanism for allocating an object, when complimentary investment is possible. Also, we propose a mechanism that allows to implement the desired outcome. At least in theory, the ε -coupon auction proposed herein has several important advantages under this environment. First, it has a unique robust equilibrium. Second, this equilibrium yields efficient allocation with near certainty (see Theorem 4). An ε -coupon auction seems simple and intuitive enough to have viable practical applications. Albeit, no amount of theorizing can guarantee that it performs well with human decision makers. Thus, comparison of an ε -coupon auction and other types of auctions may be a high pay-off project for an experimental economist.

An ε -coupon auction might be preferable to a Vickrey auction even if it is not certain whether or not the environment allows for both exogenous shocks and endogenous actions. In the extreme

case where no information revelation takes place in the signalling round, an ε -coupon auction yields a negligibly small loss in efficiency relative to a one-round second-price auction. However, as long as information revelation occurs in the signalling round of the ε -coupon auction, the additional information is likely to improve performance of the auction conducted in the second round. With independent private values, ε -coupon auction is efficient, but, obviously, using a simpler efficient mechanism is more practical in this case. However, it is reassuring that using an ε -coupon auction does no harm even if the environment has no features we explore. In short, ε -coupon auction seem to offer substantial benefits with a minimum downside.

References

- Athey, S. (2001) "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information", *Econometrica*, **69**, 861-890.
- Ausubel, L. (1998) "An Efficient Ascending-Bid Auction for Multiple Objects", mimeo.
- Bergemann, D. and Valimaki, J. (2000) "Information acquisition and efficient mechanism design" Cowles Foundation Discussion Paper 1248.
- Bergemann, D. and Valimaki, J. (2002) "Information acquisition and efficient mechanism design" *Econometrica*, 70(3):1007-1033.
- Bulow, J., Huang, M. and Klemperer, P. (1999) "Toeholds and Takeovers", *Journal of Political Economy*, **107**, 427-454.
- Cho, I-K. and Kreps, D. (1987) "Signalling Games and Stable Equilibria", *Quarterly Journal of Economics*, **102**, 179-222.
- Dasgupta, P. and Maskin, E. (2000) "Efficient Auctions", *Quarterly Journal of Economics*, **95**, 341-388.
- Daniel, K. and Hirshleifer, D. (1995) "A Theory of Costly Sequential Bidding", mimeo.
- Eso, P. and Szentes, B. (2003) "The one who controls the information appropriates its rents", mimeo.
- Fishman, M. (1988) "A Theory of Pre-emptive Takeover Bidding", *RAND Journal of Economics*, **19**, 88-101.
- Fullerton, R. and McAfee, P. (1999) "Auctioning Entry into Tournaments", *Journal of Political Economy*, **107**, 573-605.
- Jehiel, P. and Moldovanu, B. (1998) "Efficient Mechanism Design: The Case of Interdependent Valuations", mimeo.
- Haile, P. (2000) "Partial Pooling at the Reserve Price in Auctions with Resale Opportunities," *Games and Economic Behavior*, 33(2), 231-248.
- Haile, P. (2001) "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales," *American Economic Review*, 92(3), 399-427.

- Haile, P. (2003) "Auctions with Resale," *Journal of Economic Theory*, forthcoming
- Klemperer, P. (1999), "Auction Theory: A Guide to the Literature", *Journal of Economic Surveys*, **13**, 227-286.
- Klemperer, P., ed. (2000) *The Economic Theory of Auctions*, (Edward Elgar).
- Kohlberg, E. and Mertens, J. (1986) "On the Strategic Stability of Equilibria", *Econometrica*, **54**, 1003-1038.
- Krishna, V. (2002) *Auction Theory*, Academic Press, New York.
- Krishna, V. and Perry, M. (1998) "Efficient Mechanism Design", mimeo.
- Levin, D., and Smith, J.L. (1994) "Equilibrium in Auctions with Entry", *American Economic Review*, **84**, 585-599.
- Levy, H. (1992) "Stochastic Dominance And Expected Utility: Survey and Analysis", *Management Science*, **38**, 555-591.
- Lixin Ye (2000) "Indicative Bidding", mimeo.
- Maskin, E. (1992) "Auctions and Privatization", in Siebert (ed.), *Privatization*, 115-36.
- Maskin, E. (2000) "Auctions, Development, and Privatization: Efficient Auctions with Liquidity-constrained buyers", *European Economic Review*, **44**, 667-681.
- Maskin, E. and Sjostrom, T. (1999) "Implementation Theory", mimeo.
- McAfee, P. and McMillan, J. (1987), "Auctions with Entry", *Economic Letters*, **23**, 343-347.
- McAfee, P., Quan, D., and Vincent, D. (2002) "How to Set Minimum Acceptable Bids, with Application to Real Estate Auctions", *Journal of Industrial Economics*, forthcoming
- McAfee, P., Takacs, W. and Vincent, D. (1999) "Tarrifying Auctions", *Rand Journal of Economics*, **30**, no. 1.
- Milgrom, P. (1981) "Rational Expectations, Information Acquisition, and Competitive Bidding", *Econometrica*, **49**, 921-943.
- Milgrom, P. (2004) "Putting Auction Theory to Work, mimeo.
- Milgrom, P. and Weber, R. (1982) "A Theory of Auctions and Competitive Bidding", *Econometrica*, **50**, 1089-1122.
- Milgrom, P. and Shannon, C. (1994) "Monotone Comparative Statics", *Econometrica*, **62**, 157-180.
- Myerson, R. (1979) "Incentive Compatibility and the Bargaining Problem", *Econometrica*, **47**, 61-73.
- Myerson, R. and Satterthwaite, M.A. (1983) "Efficient Mechanism for Bilateral Trade", *Journal of Economic Theory*, **29**, 265-281.
- Perry, M. and Reny, P. (1999) "An Ex-Post Efficient Auction", mimeo.
- Riordan, M.H. and Sappington, D. (1987) "Awarding Monopoly Franchises", *American Economic Review*, **77**, 375-387.
- Vickrey, W. (1961) "Counterspeculation, Auctions, and Competitive Sealed Tenders", *Journal of Finance*, **16**, 8-37.

Appendix

For any number (function) x , let $x^+ = \max\{x, 0\}$. A random variable X (first-order) *stochastically dominates* a random variable Y (denoted $X \succeq Y$) if and only if for cumulative density functions, one has $F_X(z) \leq F_Y(z)$ for any $z \in \mathbf{R}$. An equivalent condition is that $E[h(X)] \geq E[h(Y)]$ for any increasing function h (e.g., Levy, 1992).

Lemma A. (1) *Suppose that X, Z and Y, W are random variables, and in both pairs variables are independent of each other. Suppose that $X \succeq Y, W \succeq Z$. Then $X - Z \succeq Y - W$.*

(2) *For any random variables X and Y such that $X \succeq Y$, and a random variable Z , which is independent of X, Y , $\max\{X, Z\} \succeq \max\{Y, Z\}$.*

(3) *For any random variables X and Y such that X stochastically dominates Y , and any constant $z \geq 0$,*

$$E[X + z]^+ - E[X]^+ \geq E[Y + z]^+ - E[Y]^+.$$

(4) *For any independent random variables X, Y, Z such that $X \succeq Y$, and any constant $t \geq 0$, $E[\max\{X + t, Y, Z\}] \succeq E[\max\{X, Y + t, Z\}]$.*

(5) *Let $q(x, y)$ be a continuous function increasing in both arguments, and let X, Y be two random variables. For any realizations $x_1 > x_2$ of the random variable X , the distribution of Y conditional on x_1 (first-order) stochastically dominates the distribution of Y conditional on x_2 . Then $q(x_1, Y) \succeq q(x_2, Y)$.*

Proof. (1)-(3) are straightforward calculations.

(4) For any numbers x, y , and z , $\max\{x, y\} = (x - y)^+ + y$. We start with the following identities

$$\begin{aligned} \max\{X + t, Y, Z\} &= (X + t - \max\{Y, Z\})^+ + \max\{Y, Z\}, \\ \max\{X, Y, Z\} &= (X - \max\{Y, Z\})^+ + \max\{Y, Z\}. \end{aligned}$$

Then

$$\begin{aligned} \max\{X + t, Y, Z\} - \max\{X, Y, Z\} &= (X + t - \max\{Y, Z\})^+ - (X - \max\{Y, Z\})^+, \\ \max\{X, Y + t, Z\} - \max\{X, Y, Z\} &= (Y + t - \max\{X, Z\})^+ - (Y - \max\{X, Z\})^+. \end{aligned}$$

From (2), we know that $\max\{X, Z\} \succeq \max\{Y, Z\}$. (1) implies that $X - \max\{Y, Z\} \succeq Y - \max\{X, Z\}$. Using (3) completes the proof.

(5) Define $\tau(x, z)$ to satisfy $q(x, \tau(x, z)) = z$. Clearly, $\tau(x, z)$ is increasing in z . Now $F_{q(x_1, Y)}(z) = F_{Y|x_1}(\tau(x_1, z)) \leq F_{Y|x_1}(\tau(x_2, z))$ and $F_{q(x_2, Y)}(z) = F_{Y|x_2}(\tau(x_2, z)) \geq F_{Y|x_1}(\tau(x_2, z))$, the latter inequality following from the fact that $Y|x_1 \succeq Y|x_2$. Therefore, for any z , $F_{q(x_1, Y)}(z) \leq F_{q(x_2, Y)}(z)$. ■

Proof of Theorem 1. (1) Let $\mathbf{a}(m) = \mathbf{a}(m, \mathbf{s})$ denote the vector of actions, where the agents with the highest m first-period signals invest, while the other $N - m$ agents skip the possibility.

To prove Theorem 1, we need to establish the following claim: *Consider vectors of actions a and a' such that $\sum_i a_i = \sum_i a'_i$, $a_i = a'_i$ for all $i \neq j, k$, and let $a_j = 1$, $a_k = 0$, $a'_j = 0$, and $a'_k = 1$. If $s_j \geq s_k$, then the expected social surplus from a is greater than that from a' .* Indeed, Let $\tilde{\mathbf{a}}$ be a vector of actions with $\tilde{a}_j = \tilde{a}_k = 0$ and $\tilde{a}_i = a_i = a'_i$ for all $i \neq j, k$. Lemma A (4) yields that $V(s_j, T_j) \succeq V(s_k, T_k)$ whenever $s_j \geq s_k$. Now one can use Lemma A (3) with the constant $b\tilde{a}_j = b\tilde{a}_k$, and the claim is proved.

This shows that a vector of actions maximizing the expected social surplus must be of the form $\mathbf{a}(m)$ for some m , $0 \leq m \leq N$. Since there is a finite number of possible m 's, there exists some r^* such that $\mathbf{a}(r^*)$ is the global maximizer of the expected social surplus. This completes the proof of Theorem 1.

(2) Definition of $G_i(\mathbf{s}, \mathbf{a}_{-i})$ implies that an action vector maximizing the social surplus must satisfy $G_i(\mathbf{s}, \mathbf{a}_{-i}) \geq 0$ when $a_i = 1$ and $G_i(\mathbf{s}, \mathbf{a}_{-i}) \leq 0$ when $a_i = 0$. We introduce a function $g_i(\mathbf{s}, \mathbf{a}_{-i})$ defined as the change in the expected utility of the agent i as a result of taking an action instead of skipping it, and prove the following claim:

$$g_i(\mathbf{s}, \mathbf{a}_{-i}) = G_i(\mathbf{s}, \mathbf{a}_{-i}).$$

Let $Z = \max_{j \neq i} \{V(s_j, t_j) + ba_j\}$, and $X = V(s_i, t_i)$. By definition,

$$g_i(\mathbf{s}, \mathbf{a}_{-i}) = E[X + b - Z]^+ - E[X - Z]^+.$$

Using the formula $\max\{x, y\} = (x - y)^+ + y$, we get

$$\begin{aligned} G_i(\mathbf{s}, \mathbf{a}_{-i}) &= E[\max\{X + b, Z\}] - E[\max\{X, Z\}] \\ &= E[X + b - Z]^+ + E[Z] - (E[X - Z]^+ + E[Z]) \\ &= E[X + b - Z]^+ - E[X - Z]^+ = g_i(\mathbf{s}, \mathbf{a}_{-i}), \end{aligned}$$

as claimed.

Now observe that if \mathbf{a} is a solution to the social planner's problem, then $G_i(\mathbf{s}, \mathbf{a}_{-i}) \geq 0$ when $a_i = 1$ and $G_i(\mathbf{s}, \mathbf{a}_{-i}) \leq 0$ when $a_i = 0$. Indeed, if $G_i(\mathbf{s}, \mathbf{a}_{-i}) < 0$ when $a_i = 1$, the agent's i switch from acting to non-acting would strictly increase the expected social surplus, contradicting the choice of \mathbf{a} . Similarly, $G_i(\mathbf{s}, \mathbf{a}_{-i}) \leq 0$ when $a_i = 0$. Then the claim proved above asserts that we have $g_i(\mathbf{s}, \mathbf{a}_{-i}) \geq 0$ for agents that invest, and $g_i(\mathbf{s}, \mathbf{a}_{-i}) \leq 0$ for others. Thus, no agent has incentives to deviate, and (ii) of Theorem 1 is proved. ■