Reprinted from JOURNAL OF ECONOMIC THEORY
All Rights Reserved by Academic Press, New York and London
Vol. 44, No. 1, February 1988
Printed in Belgium

Search Mechanisms*

R. PRESTON MCAFEE

Department of Economics, University of Western Ontario, London, Ontario, N6A 5C2, Canada

AND

JOHN MCMILLAN

Department of Economics, University of Western Ontario, London, Ontario N6A 5C2, Canada, and Graduate School of International Relations and Pacific Studies, University of California, San Diego, La Jolla, California 92093

Received August 8, 1986; revised January 9, 1987

Extending the Revelation Principle to a case in which it is costly for the principal to communicate with any agent, we show that there is a sequential direct mechanism that is optimal in the class of all mechanisms. We then apply this result to the problem of a monopsonist seeking to buy an indivisible good from one of a set of possible sellers with unobservable production costs. With costly communication, the monopsonist's optimal procurement mechanism is a combination of reservation-price search and auction. *Journal of Economic Literature* Classification Numbers: 022, 026, 213. © 1988 Academic Press, Inc.

1. INTRODUCTION

The traditional approach to modelling the optimizing behavior of a monopolist or monopsonist is to take as given the selling or buying policy—such as posting a fixed price, or choosing among various types of auction, or searching sequentially—and to optimize within this given institution. Thus one solves for the best price to post, or the best of the given auction forms to use, or the best reservation price. A more fundamental approach is to optimize over institutions, without constraining the allowable types of selling or buying policies. When should a monopolist

* We thank Steven Lippman, Peter Morgan, Michael Peters, Michael Riordan, Michael Rothschild, Ariel Rubinstein, Joel Sobel, and Associative Editor and a referee for helpful comments.

0022-0531/88 $3.00
Copyright © 1988 by Academic Press, Inc.
All rights of reproduction in any form reserved.

choose to sell by auction; when should he post a fixed price? When is sequential search the best of all possible buying policies? Optimization over institutions is made possible by the application of the Revelation Principle.

Consider a monopsonist who wishes to acquire one unit of a good from one of a set of possible sellers. The buyer has the ability to commit himself in advance to his buying policies: thus he is endowed with considerable bargaining power. His power to extract surplus is limited, however, by his inability to observe the potential sellers' production costs. In addition he must incur a search cost in communicating with any potential seller.

The monopsonist designs his optimal buying mechanism. It is now well known that, in the absence of communication costs, the mechanism that is optimal for the monopsonist is a sealed-bid or oral auction, augmented by a reserve price.[1] This paper will show that, with costly communication, the optimal mechanism has the buyer approaching the potential sellers in sequence: the optimal mechanism works like a marriage of sequential search and auction.

Most existing search models have many buyers and many sellers.[2] The assumption of a single buyer with commitment ability does not fit the usual interpretation of search theory as representing people buying goods for their own consumption. It might, however, represent a large industrial buyer procuring inputs from other firms, or a government contracting out to the private sector the production of a public good. It might also represent the hiring process of a monopsonistic employer. There are several ways a large buyer might achieve commitment. For instance, in the case of government contracting, the government official responsible for the decision is required to follow procedures that are explicitly and precisely set out in a publicly available book of rules. Alternatively, reputational effects might produce commitment: the cost to the monopsonist of reneging on his announced policy might be the inability to credibly commit himself in the future and the consequent loss of future bargaining power.

In most equilibrium search models, the sellers are Stackelberg leaders and the buyers are Stackelberg followers: in other words, the sellers are assumed to be able to commit themselves to their price offers. The present model reverses the usual commitment assumption: the buyer leads and the sellers follow.[3]

The search cost can be given its usual interpretation as the cost of locating and contacting potential sellers. Alternatively, in the contracting interpretation, it might be the cost of checking that a potential supplier is

[1] Myerson [15], Riley and Samuelson [18], McAfee and McMillan [11].

[2] Exceptions to this statement, search models with small numbers of agents, include Carlson and McAfee [2] and Reinganum [17].

[3] In the model of Wilde [22], the buyers and the sellers move simultaneously.

capable of doing the work for which he has bid. Or it may be that there is competition over design as well as price, so that the bids differ in several quality dimensions as well as in price. The buyer must reduce these multidimensional characteristics to a single-dimensional comparison in order to decide which is the best offer. This may be time-consuming and therefore costly to the buyer.[4] In the job-market interpretation, the search cost might represent the cost to the employer of checking the credentials of a potential employee.

In the usual formulation of the Revelation Principle, information, though asymmetric, can be transmitted without cost. In Section 3 we extend the Revelation Principle in the generalized principal-agent framework of Myerson [16], with both adverse selection and moral hazard, to a case where communication is costly for the principal. We show that a *sequential direct mechanism* is optimal in the class of all mechanisms. In a sequential direct mechanism, the principal asks the agents their types (and nothing else) in sequence. At any time the principal may stop asking agents their types, so that he need not communicate with all of the agents. In addition, it can be assumed without loss of optimality that those agents who are asked reveal their types truthfully and execute the decision the principal recommends for them. Thus, one can construct an optimal mechanism by optimizing over the class of sequential direct mechanisms subject to the usual incentive-compatibility constraints.

The extended Revelation Principle is then applied in Section 4 to the monopsonist's search problem. The main result is that the monopsonist sets two cut-off cost levels x_0 and x^*, with $x_0 > x^*$, and proceeds sequentially. If he finds a seller with cost less than x^*, he immediately buys from that seller; otherwise he continues searching. If he exhausts the entire set of potential sellers, he buys from the lowest-cost seller, provided his cost is no higher than x_0 (where x_0 is determined by the buyer's fallback option—in-house production, say). The price the monopsonist must pay is the expected second-lowest cost.

In the limiting case of infinitely many potential sellers, the optimal mechanism is pure reservation-price search.[5] When the cost of communication goes to zero, the mechanism reduces to the usual optimal auction. Thus reservation-price search and auctions are inherently related, in that each emerges as a special case of the same mechanism.

In the job-market interpretation, the optimal mechanism works as follows. The employer, at some cost to himself, evaluates the credentials of

[4] These costs can be large in practice. For example, Fox [4, p. 269] cited a U.S. Department of Defense contract in which government personnel spent 182,000 man-hours evaluating proposals from four prospective contractors.

[5] Riley and Zeckhauser [19] previously obtained this limit result.

the job applicants one at a time. Suppose the applicants' abilities net of their opportunity costs can be measured on a single-dimensional scale. If an applicant's qualifications are evaluated by the employer as being inferior to the cut-off level x_0, he is immediately rejected. If an applicant is judged to be better than the cut-off level x^*, he is immediately hired. If his qualifications put him between x_0 and x^*, he is told "don't call us, we'll call you." Then, if no applicant who is better than x^* has been found by the time all potential employees have been interviewed, he is offered the job if he has the best available qualifications.

In the case of government procurement, the government agency faces a "make-or-buy" decision: our theorem has the common-sense implication that the buyer rationally produces the item in-house if the search cost is relatively high and the cost of in-house production is relatively low. One controversial aspect of U.S. military procurement is that a majority of contracts are let on a sole-source basis (Fox [4]): the theorem shows that there are combinations of the search cost and the distribution of bidders' production costs such that it is rational to solicit only one bid. Otherwise, according to the theorem, bids should be solicited in sequence. Also, there are circumstances under which minimizing procurement costs requires the government to produce the item in-house even after having found a firm which could produce it with a lower production cost (because the price it would have to pay the firm is higher than its in-house production cost).

Note that the model of this paper is an equilibrium search model (although the notion of equilibrium is not the usual search equilibrium[6]). The buyer's and the sellers' optimization give rise to a dispersion of price offers in equilibrium. Why does the usual argument (due to Diamond [3]; see also Rothschild [20]), that search costs result in all sellers charging the monopoly price no matter how many sellers there are, not apply here? In the Diamond argument, any price less than the monopoly price cannot persist in equilibrium because a seller known that, once the buyer has incurred the search cost and has received his price quotation, he can slightly raise his price and it remains in the buyer's interest to accept the higher price rather then incur further search costs. The argument does not apply here because we have endowed the buyer with commitment ability. The seller does not raise his price above the reservation price because he knows that the buyer has irrevocably committed himself not to pay more, even though it may be in his *ex post* interest to pay more. Thus the buyer's commitment eliminates the Diamond monopoly-price equilibrium.

[6] On equilibrium price dispersion, see Burdett and Judd [1], Carlson and McAfee [2] and the references therein.

2. Searching for the Lowest Bid

A monopsonist wishes to buy one unit of an indivisible good. The buyer is assumed to be able to commit himself to a purchasing policy. There are n potential sellers, who vary in that they may have different production costs, which only they themselves know (or, in the job market interpretation, different potential employees have different abilities net of oppotunity costs). Denote a seller's production cost by x, and suppose that costs are identically and independently distributed[7] as $F(x)$, with $F'(x) = f(x)$ and $F \in C^1$, $F(0) = 0$. The buyer is able to produce the good himself at a commonly known cost of $z_0 > 0$. (The case in which the buyer has no such option is represented by $z_0 = \infty$.) Both the potential sellers and the buyer are risk neutral.

The buyer incurs a cost $c > 0$ every time he contacts a potential seller. Assume that he cannot avoid incurring the search costs by, for example, publicly advertising the price he would be willing to pay.

Define a function J by

$$J(x) = x + \frac{F(x)}{f(x)}, \tag{1}$$

and assume J is strictly increasing on $\{x \mid 0 < F(x) < 1\}$. (This is the analogue—for buying as opposed to selling—of the J function of Maskin and Riley [9] and the c_i function of Myerson [15]: it is assumed to be strictly increasing for the same reason as in those papers.) Define $x_0 = J^{-1}(z_0)$. Since the function J will figure prominently in the analysis that follows, the following lemma is useful as an aid to understanding.

Lemma 1. *Let s^i represent the* ith *order statistic of the set of seller's costs. Then*

$$E[J(s^1)] = E[s^2]. \tag{2}$$

Proof. The density of the second order statistic is

$$n(n-1)[1 - F(x)]^{n-2} F(x) f(x).$$

[7] This is an independent-private-values model, in the terminology of Milgrom and Weber [12].

Thus the expected value of the second order statistic is

$$\begin{aligned} Es^2 &= \int_0^\infty xn(n-1)[1-F(x)]^{n-2}F(x)f(x)\,dx \\ &= -n\int_0^\infty [xF(x)]\frac{d}{dx}[1-F(x)]^{n-1}\,dx \\ &= -nxF(x)[1-F(x)]^{n-1}\big|_0^\infty + n\int_0^\infty [xf(x)+F(x)][1-F(x)]^{n-1}\,dx \\ &= \int_0^\infty J(x)\,n[1-F(x)]^{n-1}f(x)\,dx. \end{aligned} \tag{3}$$

Thus, since $n[1-F(x)]^{n-1}f(x)$ is the density of the first order statistic, (2) holds. Q.E.D.

It follows from Lemma 1 that the expected difference between the lowest-cost seller's cost and the second-lowest-cost seller's cost is the expected value of $F(x)/f(x)$. Thus, by the usual auction-theory intuition (see McAfee and McMillan [11]), the winning bidder's expected profit is the expected value of $F(x)/f(x)$ and his expected payment is the expected value of $J(x)$. $F(x)/f(x)$ is therefore to be interpreted as the expected informational cost borne by the principal, resulting from the sellers' private information. Note, however, that the amount paid to the successful seller is not $J(x)$. The seller has an extra piece of information: he knows his own cost. The seller will quote a price equal to his expectation of the next-lowest cost, conditional on the level of his own cost: this price will be derived in Theorem 9 (Eq. (27)). Only in expectation does this price equal $J(x)$.

To be applicable to this problem, the Revelation Principle must be generalized to admit communication costs. This is done, in a model which is more general than the procurement problem just stated, in the next section. Much of the terminology, notation, and method of analysis used in Section 3 is borrowed from Myerson [16]. (The reader prepared to accept that the Revelation Principle, suitably modified into a sequential form, does apply to a model with communication costs can skip the next section and go directly to Section 4.)

3. The Revelation Principle with Costly Communication

There are $n < \infty$ agents, indexed by $i \in N = \{1, 2, ..., n\}$, who behave non-cooperatively. Denote by Ω the set of ordered subsets of agents: $\Omega = \{(i_1, ..., i_k) \mid k \leqslant n,\ 1 \leqslant i_j \leqslant n, \text{ and } j < m \leqslant k \Rightarrow i_j \neq i_m\}$. Agent i has type $t_i \in T_i$, which only he can observe. Let $T = X_{i \in N} T_i$. There is a probability distribution $P: T \rightarrow [0, 1]$ which is common knowledge, and we assume that, given his type t_i, agent i uses the Bayesian posterior of P as the probability of the vector types. Agent i can make decision $d_i \in D_i$; the principal cannot directly control the agent's decision.[8] The principal makes a decision $d_0 \in D_0$. Let $D = X_{i=0}^{n} D_i$. The agents have von Neumann–Morgenstern utility functions $U_i: D \times T \rightarrow R$.

Since communicating with an agent is costly for the principal, the principal's utility depends upon his communications. We shall assume that his utility depends only on the set of agents actually communicated with and on $(d, t) \in D \times T$. Suppose that, once the principal has communicated with agent i, there are no additional costs incurred for further communication with agent i. (It will be shown that this restriction is without loss of generality.) Let $U_0: D \times T \times \Omega \rightarrow R$ be the principal's utility function. Thus, we are in essence assuming fixed costs of communication. Although communication costs vary with types and decisions, they are invariant to the actual message sent.[9]

We seek to show that Revelation Principle extends to the case of costly communication. Why might it not hold; why might truth-telling not be an equilibrium? Essentially, the Revelation Principle must apply if there is nothing the principal can do to avoid bearing the communication cost. We ensure this by placing two restrictions on the nature of the communication.

We firstly restrict attention to what we shall call *principal-centered* mechanisms; by this is meant that the only type of communication that takes place is agent-to-principal of principal-to-agent. If, on the contrary, agents were allowed to communicate among themselves, it would be in the principal's interest to incur the cost of communicating with only one agent, and to have that agent learn the types of the other agents. (It appears to be intractable to make agent-to-agent communication possible but costly.) Note that, in the absence of communication costs as in the usual formulation of the Revelation Principle, allowing only communication between principal and agent is not a restriction, because the principal could commit to passing signals from agent i to agent j without himself observing the signal (that is, the principal would commit himself not to condition any of his future signals or decisions on this communication). Thus there is a sense in which the restriction to principal-centered mechanisms is no less general than the usual Revelation Principle analysis with costless com-

[8] In the search problem stated in the last section, there is no decision for the agents to take. Nevertheless, in this section's extension of the Revelation Principle we allow for the possibility of decisions by the agents for the sake of generality and comparability with Myerson [16]. Such decisions could be added to the model of Section 2 by giving the selected seller some ex post control over the cost he incurs, as in McAfee and McMillan [10].

[9] For analyses of the case in which the cost of communication does depend on the size of the signal sent, see E. Green [5] and J. Green and Laffont [6].

munication, in that the usual analysis can be interpreted as being principal-centered. Note also that there is a relationship between the assumption of no agent-to-agent communication and the assumption that the agents behave noncooperatively, in that communication is a prerequisite for cooperation in a static game: in this sense, the absence of agent-to-agent communication is a sufficient condition for noncooperative behavior.

We also restrict attention to what we shall call *principal-initiated* mechanisms. This is to avoid the possibility that the principal might use the mechanism itself as a communication device and in so doing avoid incurring the costs of communication. For example, suppose the principal wishes to buy a good and chooses a mechanism that dictates that he will accept the first bid of not more than x^* that he receives from an agent. Then any agent who values the good at less than x^*, knowing that is the principal's decision rule, will submit a bid of x^*. The principal then accepts the first bid he receives and shuts off communication, incurring the cost of only one communication. In this example, the mechanism itself communicates the value x^* costlessly to all agents. In a principal-initiated mechanism, this cannot happen. A mechanism is principal-initiated if no communication occurs between the principal and a particular agent unless the principal first sends a message to the agent; agents never initiate communication. (For example, one may imagine that the principal must contact the agent, explain the mechanism to him, and demonstrate that he is bound to it.) We assume, then, that if the principal does not communicate with agent i, then agent i chooses a decision $\hat{d}_i(t_i)$, which is the same regardless of which mechanism the principal has choosen.

To the extent that communication takes time, the principal and the agents are assumed not to discount future returns (or, alternatively, the communication process is sufficiently fast that discounting can be ignored as an approximation). Thus simultaneous communication can be ignored without loss of generality. The principal has the option of mimicking a simultaneous mechanism by receiving the communications sequentially but committing himself to make no use of the information until all of the information is received; that is, until the "simultaneous" information is fully received, no action or signal of the principal is conditioned on the initial information. Hence in what follows, attention will be restricted to sequential mechanisms. At some increased notational complexity, real time could be included in the communication model. This would in general lead to simultaneous signals, as in the search model of Morgan [13] and Morgan and Manning [14]. The extension to this case is straightforward if it is feasible for the principal to delay sending signals. Indeed, the proof of the Revelation Principle is consistent with this case.

A principal-initiated, principal-centered mechanism allows an exchange of signals between the principal and some or all of the agents and, when this exchange of information has ended, culminates in a vector of decisions $(d_0, d_1, ..., d_n) \in D$.

In a *direct* mechanism, each agent simply reports his type from T_i to the principal, who responds by suggesting a decision from D_i for the agent. In a *sequential direct* mechanism, the principal communicates with the agents one at a time. An agent is *honest* if he correctly reports his type when asked. An agent is *obedient* if he takes the decision recommended by the principal. The sequential direct mechanism is incentive compatible when honest and obedient strategies form a Bayes–Nash equilibrium.

We now show how the Revelation Principle extends to the case of costly communication. A direct sequential mechanism is optimal in the class of all principal-initiated mechanisms.

LEMMA 2. *Corresponding to any equilibrium σ for any principal-initiated mechanism μ, there is a direct sequential incentive-compatible mechanism μ^* in which, for each vector of types $t \in T$, an honest and obedient strategy σ^* produces the same distribution of decisions and agents communicated with (d, ω), $d \in D$, $\omega \in \Omega$.*

The result is proven in the Appendix, following Myerson [16].

The foregoing analysis assumed that the principal incurred a cost only upon the first communication with any particular agent. However, Lemma 2 showed that the principal need only communicate at most twice with each agent; either not all, or once to ask his type and once to suggest his decision. Clearly, therefore, if the second and subsequent communications with any one agent are costly to the principal, the principal optimizes by using a direct sequential mechanism. The assumption that only the first communication with any agent is costly for the principal is therefore without loss of generality.

4. THE MONOPSONIST'S OPTIMAL MECHANISM

We return now to the problem of procurement with search costs stated in Section 2. From Lemma 2, we know that there is an optimal direct mechanism that has the buyer sequentially asking the potential sellers their types (production costs), x. Let the subscript i denote the ith potential seller asked. Let $x_1, ..., x_n$ be the random variables that are their responses. Let $y_k = \min\{x_0, x_1, ..., x_k\}$ be the lowest of the first k responses, together with x_0.

At the kth stage, the buyer chooses between producing the good himself at cost z_0; buying the good from seller i, $i = 1, ..., k$; or continuing to ask further potential sellers. Given $x_0, x_1, ..., x_{k-1}$, denote by γ_k^0, γ_k^i, $i = 1, ..., k$, and α_k the set of kth responses to x_k in which he makes these respective decisions. Thus, if $x_k \in \gamma_k^0(x_0, x_1, ..., x_{k-1})$, the buyer decides after the kth stage to produce the good himself; if $x_k \in \gamma_k^i(x_0, x_1, ..., x_{k-1})$, he purchases it from seller i; and if $x_k \in \alpha_k(x_0, x_1, ..., x_{k-1})$, he takes his $(k+1)$th observation. Let

$$\Gamma_k^i = \{(x_1, ..., x_k) \mid x_k \in \gamma_k^i(x_0, ..., x_{k-1})\}, \qquad i = 0, ..., k, \tag{4}$$

and

$$\beta_k^i(z) = \{(x_1, ..., x_{i-1}, x_{i+1}, ..., x_k) \mid (x_1, ..., x_{i-1}, z, x_{i+1}, ..., x_k) \in \Gamma_k^i\}, \qquad i = 0, 1, ..., k. \tag{5}$$

Thus Γ_k^i is the set of others' responses such that bidder i wins in round k if he reports x_i; $\beta_k^i(z)$ is the set of others' responses such that i wins in the kth round if he reports z. The arguments of α_k and γ_k^i will be suppressed for brevity.

It follows that, if a seller reports a cost of z, his probability of winning the contract is

$$\mu(z) = \sum_{k=1}^{n} \sum_{i=1}^{k} \int_{\beta_k^i(z)} f_k(x_{-i})\, dx_{-i}, \tag{6}$$

where $x_{-i} = (x_1, ..., x_{i-1}, x_{i+1}, ..., x_k)$ and $f_k(x_{-i}) = \prod_{j=1, j \neq i}^{k} f(x_j)$.

It follows from Lemma 2 that there is a function $A_k^i(x)$ that represents the amount i is paid, given that he is asked to supply the good in the kth round. Assume that i is paid if and only if he is asked to supply the good: this is without loss of generality because of risk neutrality. Denote by $p(z)$ a sellers' average payment given that he is asked to supply the good and that his reported cost is z; $p(z)$ satisfies

$$\mu(z)\, p(z) = \sum_{k=1}^{n} \sum_{i=1}^{k} \int_{\beta_k^i(z)} A_k^i(x_{-i}, z)\, f_k(x_{-i})\, dx_{-i}. \tag{7}$$

If a seller's true cost is x and he reports z, his expected profit is

$$\pi(z) = \mu(z)[p(z) - x]. \tag{8}$$

The incentive-compatibility constraint requires that $\pi(z)$ is maximized at $z = x$. This requires

$$\frac{d}{dz}[p(z)\ \mu(z)]\bigg|_{z=x} = x\mu'(z)|_{z=x} \tag{9}$$

or

$$p(x)\, \mu(x) = -\int_x^{x_m} z\mu'(z)\, dz$$

$$= x\mu(x) + \int_x^{x_m} \mu(z)\, dz, \tag{10}$$

where x_m is defined by $p(x_m)\,\mu(x_m)=0$. Since $p(x)\geqslant x$ for $\mu(x)>0$ by the assumption of free exit, x_m satisfies $x_m=\inf\{x \mid \mu(x)=0\}$. In addition, the second-order condition requires $\mu'(x)\leqslant 0$, which is assumed to hold, and will hold in the solution.

LEMMA 3. *The expected payment to the successful bidder is*

$$\tau=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\Gamma_k^i} J(x_i)\,f(x)\,dx. \qquad (11)$$

Proof.

$$\tau=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\Gamma_k^i} A_k^i(x)\,f(x)\,dx \qquad \text{(by definition of } A_k^i\text{)}$$

$$=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_0^{\infty}\left[\int_{\beta_k^i(z)} A_k^i(x_{-i},z)\,f_k(x_{-i})\,dx_{-i}\right] f(z)\,dz$$

$$=\int_0^{\infty}\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\beta_k^i(z)} A_k^i(x_{-i},z)\,f_k(x_{-i})\,dx_{-i}f(z)\,dz$$

$$=\int_0^{\infty} p(z)\,\mu(z)\,f(z)\,dz \qquad \text{(from (7))}$$

$$=\int_0^{x_m} p(z)\,\mu(z)\,f(z)\,dz$$

$$=\int_0^{x_m} zf(z)\,\mu(z)\,dz+F(x)\int_z^{x_m}\mu(x)\,dx\Big|_0^{x_m}+\int_0^{x_m}F(z)\,\mu(z)\,dz \qquad \text{(from 10))}$$

$$=\int_0^{\infty} J(z)\,f(z)\,\mu(z)\,dz \qquad \text{(from (1))}$$

$$=\int_0^{\infty} J(z)\,f(z)\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\beta_k^i(z)} f_k(x_{-i})\,dx_{-i}\,dz \qquad \text{(from (6))}$$

$$=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_0^{\infty}\int_{\beta_k^i(z)} J(z)\,f(z)\,f_k(x_{-i})\,dx_{-i}\,dz$$

$$=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\Gamma_k^i} J(x_i)\,f(x_i)\,f_k(x_{-i})\,dx_{-i}\,dx_i$$

$$=\sum_{k=1}^{n}\sum_{i=1}^{k}\int_{\Gamma_k^i} J(x)\,f(x)\,dx. \qquad \text{Q.E.D.}$$

Define some more notation: let

$$h_k = c + \sum_{i=0}^{k} \int_{\gamma_k^i(x_0, x_1, ..., x_{k-1})} J(x_i)\, f(x_k)\, dx_k; \tag{12}$$

$$\phi_k = h_k(x_0, x_1, ..., x_{k-1}) + \int_{\alpha_k(x_0, x_1, ..., x_{k-1})} f(x_k)\, \phi_{k+1}(x_0, x_1, ..., x_k)\, dx_k; \tag{13}$$

$$\phi_n = h_n. \tag{14}$$

The interpretation of these variables is as follows. Suppose the buyer has observed $x_1, ..., x_{k-1}$. His expected cost if he takes exactly one more observation is h_k since, when he accepts a type x_i, he pays him of average $J(x_i)$, as implied by Lemmas 1 and 3. His total expected cost given that he does not stop at the $(k-1)$th observation is ϕ_k because he pays h_k if he stops at the kth observation and pays ϕ_{k+1} if he continues beyond the kth observation. Hence ϕ_k is the buyer's total expected cost associated with taking the kth observation.

LEMMA 4. *The total expected cost incurred by the buyer is, for* $1 \leqslant k \leqslant n-1$:

$$\phi = h_1 + \int_{\alpha_1} f(x_1) \Big[h_2(x_1) + \int_{\alpha_2} f(x_2) \Big[\cdots + \Big[\int_{\alpha_k} f(x_k)\, \phi_{k+1}(x_0, x_1, ..., x_k)\, dx_k \Big] \cdots \Big] dx_1. \tag{15}$$

Proof. From (11)

$$\phi = \sum_{k=1}^{n} \Big[(c_0 + kc) \int_{\Gamma_k^0} f(x)\, dx + \sum_{i=1}^{k} \int_{\Gamma_k^i} [J(x_i) + kc]\, f(x)\, dx \Big]$$

$$= \sum_{k=1}^{n} \Big[\sum_{i=0}^{k} \int_{\Gamma_k^i} [J(x_i) + kc]\, f(x)\, dx \Big]$$

$$= c + \sum_{k=1}^{n} \int_{\alpha_1} \cdots \int_{\alpha_{k-1}} \prod_{j=1}^{k-1} f(x_j) \times \Big\{ c \int_{\alpha_k} f(x_k)\, dx_k + \sum_{i=0}^{k} \int_{\gamma_k^i} J(x_i) f(x_k)\, dx_k \Big\}\, dx_{k-1} \cdots dx_n \qquad \text{(by (4))}$$

$$= h_1 + \int_{\alpha_1} f(x_1) \Big[h_2(x_1) + \int_{\alpha_2} f(x_2) \times \Big[\cdots \Big[\int_{\alpha_{n-1}} f(x_{n-1})\, h_n(x_{-n})\, dx_{n-1} \Big] \cdots \Big] dx_1 \qquad \text{(by (12))}$$

$$= h_1 + \int_{\alpha_1} f(x_1) \Bigg[h_2(x)$$

$$+ \Bigg[\cdots \Bigg[\int_{\alpha_{n-2}} f(x_{n-2})\, \phi_{n-1}(x_0, ..., x_{n-2})\, dx_{n-2} \Bigg] \cdots \Bigg] dx_1. \tag{16}$$

Equation (15) is obtained by backward induction. Q.E.D.

Clearly, the cost to the buyer of searching further, ϕ_k, depends on the reported costs $x_0, x_1, ..., x_{k-1}$. However, we now show it depends only on the minimum of these, $y_{k-1} = \min\{x_0, x_1, ..., x_{k-1}\}$.

LEMMA 5. *Minimizing the total expected cost incurred by the buyer, ϕ, implies that the cost of continuing, ϕ_k, depends only on the lowest previously observed cost, y_{k-1}.*

Proof. Recall that J is assumed to be increasing. The proof is by induction. For the base, note that, from the last line of the proof of Lemma 4 and (13), minimizing ϕ requires minimizing $\phi_n = h_n$. By (12), this occurs by putting x_k in γ_n^i when $J(x_i)$ is smallest, which occurs at $x_i = y_{n-1}$ if $x_n \geqslant y_{n-1}$, or $x_i = x_n$ if $x_n < y_{n-1}$. Thus

$$\phi_n = h_n = c + \left\{ \int_0^{y_{n-1}} J(x) f(x_n)\, dx_n + \int_{y_{n-1}}^{\infty} J(y_{n-1})\, f(x_n)\, dx_n \right.$$

$$= c + \int_0^{y_{n-1}} J(x_n) f(x_n)\, dx_n + [1 - F(y_{n-1})]\, J(y_{n-1}). \tag{17}$$

This proves the base of the induction. From (15), we must minimize ϕ_k over α_k, γ_k^i. Suppose ϕ_{k+1} depends only on $y_k = \min\{y_{k-1}, x_k\}$.

$$\phi_k = c + \sum_{i=0}^{k} \int_{\gamma_k^i} J(x_i)\, f(x_k)\, dx_k + \int_{\alpha_k} f(x_k)\, \phi_{k+1}(y_k)\, dx_k \quad \text{(by (12), (13))}$$

$$= c + \int_0^{y_{k-1}} \min\{J(x_k), \phi_{k+1}(x_k)\}\, f(x_k)\, dx_k$$

$$+ \int_{y_{k-1}}^{\infty} \min\{J(x_k), \phi_{k+1}(y_{k-1})\}\, f(x_k)\, dx_k$$

$$= c + \int_0^{y_{k-1}} \min\{J(x), \phi_{k+1}(x)\}\, f(x)\, dx + [1 - F(y_{k-1})]$$

$$\times \min\{J(y_{k-1}), \phi_{k+1}(y_{k-1})\}. \tag{18}$$

Since this depends only on y_{k-1}, the proof is complete. Q.E.D.

COROLLARY 6. *The set of reports for which, at the* kth *stage, the buyer decides to continue searching, is*

$$\alpha_k(y_{k-1}) = \{x \mid z = \min\{y_{k-1}, x\} \Rightarrow \phi_{k+1}(z) \leqslant J(z)\}. \tag{19}$$

Define

$$\psi_k(y) = \phi_k(y) - J(y). \tag{20}$$

Thus $\psi_k(y)$ is the difference between the expected cost to the buyer, at stage $k-1$, of searching further and the cost of purchasing at the current best observation.

LEMMA 7. *$\psi_k(y)$ is strictly decreasing in y.*

Proof. As the base of the induction, note that $\psi'_n = -J' < 0$ by assumption. From (20) and the last line of the proof of Lemma 5,

$$\begin{aligned}
\psi_k(y) &= c + \int_0^y \min\{0, \psi_{k+1}(x)\} f(x)\,dx \\
&\quad + \int_0^y J(x) f(x)\,dx - F(y)\,J(y) + [1 - F(y)] \min\{0, \psi_{k+1}(y)\} \\
&= c + \int_0^y [J(x) - J(y)] f(x)\,dx + \int_0^y \min\{0, \psi_{k+1}(x)\} f(x)\,dx \\
&\quad + [1 - F(y)] \min\{0, \psi_{k+1}(y)\} \\
&= c - \frac{[F(y)]^2}{f(y)} + \int_0^y \min\{0, \psi_{k+1}(x)\} f(x)\,dx \\
&\quad + [1 - F(y)] \min\{0, \psi_{k+1}(y)\},
\end{aligned} \tag{21}$$

the last line following because

$$\int_0^y J(x) f(x)\,dx = xF(x)\Big|_0^y - \int_0^y F(x)\,dx + \int_0^y F(x)\,dx = yF(y) \qquad \text{(by (1))} \tag{22}$$

so that

$$\begin{aligned}
\int_0^y [J(x) - J(y)] f(x)\,dx &= yF(y) - \left[y + \frac{F(y)}{f(y)}\right] F(y) \\
&= \frac{-[F(y)]^2}{f(y)}.
\end{aligned} \tag{23}$$

It follows from the last line of (21) that

$$\psi_k' = -J'F(y) + [1 - F(y)] \frac{d}{dy} \min\{0, \psi_{k+1}(y)\} < 0. \tag{24}$$

Q.E.D.

Define x_k^* by $\psi_{k+1}(x_k^*) = 0$. Since $\psi_{k+1}(y)$ is strictly decreasing for $0 < F(y) < 1$, there is at most one interior solution of $\psi_{k+1}(x_k^*) = 0$. From Corollary 6, x_k^* acts like a reservation price (or, more accurately, a reservation type), since the buyer continues to search if and only if $y_k > x_k^*$.

LEMMA 8. $x_1^* = x_2^* = \cdots = x_{n-1}^*$.

Proof. Suppose, by way of induction, $x_k^* = x_{k+1}^* = \cdots = x_{n-1}^*$, for some $k \leqslant n-1$. This is true for $k = n-1$. From (21)

$$\begin{aligned} \psi_k(x_k^*) &= c - \frac{[F(x_k^*)]^2}{f(x_k^*)} + [1 - F(x_k^*)] \min\{0, \psi_{k+1}(x_k^*)\} \\ &\quad + \int_0^{x_k} \min\{0, \psi_{k+1}(x)\} f(x)\, dx \\ &= \left[c - \frac{[F(x_{n-1}^*)]^2}{f(x_{n-1}^*)} \right] + \int_0^{x_k^*} \min\{0, \psi_{k+1}(x)\} f(x)\, dx \\ &= 0, \end{aligned} \tag{25}$$

since the first term in brackets is $\psi_n(x_{n-1}^*) = 0$ and $x < x_k^* \Rightarrow \psi_{k+1}(x) \geqslant \psi_{k+1}(x_k^*) = 0$. Thus $x_{k-1}^* = x_k^*$. Q.E.D.

Define x^* by

$$c = \frac{[F(x^*)]^2}{f(x^*)}. \tag{26}$$

From the proof of Lemma 8, x^* is the constant cut-off reported cost which determines whether or not the buyer continues searching. For some intuitive understanding of why this cut-off is determined by (26), recall from Lemma 1 that a seller with cost x^* makes a profit equal to the difference on average between his own cost and the cost of the second-lowest bidder, or $F(x^*)/f(x^*)$. Suppose the lowest-cost bidder the buyer has so far observed has cost of x^*. If the buyer stops searching now and buys from this bidder, the price he pays is $x^* + F(x^*)/f(x^*)$. If instead he takes one more observation and finds a lower-cost seller, he must pay the new bidder a price equal to the cost of the second-lowest-cost bidder, which is now x^*. Thus, if he searches once more and finds a lower-cost seller, he saves

$F(x^*)/f(x^*)$ on average. The probablity of finding a lower-cost seller with the next observation is $F(x^*)$. Hence the marginal expected benefit to one more observation is $[F(x^*)]^2/f(x^*)$. The marginal cost is c. Hence (26) simply equates marginal benefit to marginal cost.

The optimal mechanism can now be summarized.

THEOREM 9. *The optimal strategy for the buyer is:*

(a) *If* $x_0 = J^{-1}(z_0) < x^*$, *the buyer consults no potential sellers, and produces the good himself.*

(b) *If* $x_0 > x^*$ *and* $F(x^*) = 1$, *the buyer takes one observation and pays* $x_{\max} = \inf\{x \mid F(x) = 1\}$.

(c) *If* $x_0 > x^*$ *and* $F(x^*) < 1$, *the buyer sequentially samples the sellers until the first with a cost no greater than* x^* *is found; if he finds no such seller, he samples all of the sellers and either buys from the lowest-cost seller or produces the good himself if the lowest-cost seller's cost exceeds* x_0. *The payment to a seller with cost y is then*

$$p(y) = \begin{cases} y + [1 - F(y)]^{-(n-1)} \int_y^{x_0} [1 - F(x)]^{n-1}\, dx, & y > x^* \\ x^* + \int_{x^*}^{x_0} [1 - F(x)]^{n-1}\, dx, & y \leqslant x^*. \end{cases} \tag{27}$$

On average, the total cost to the buyer is

$$\begin{aligned} \phi = {} & \left[p(x^*) + \frac{c}{F(x^*)} \right] [1 - (1 - F(x^*))^n] + J(x_0)(1 - F(x_0))^n \\ & + \int_{x^*}^{x_0} \left[y + \frac{F(y)}{f(y)} \right] n(1 - F(y))^{n-1} f(y)\, dy \\ & - nF(x^*) \int_{x^*}^{x_0} [1 - F(x)]^{n-1}\, dx. \end{aligned} \tag{28}$$

Proof. Equations (27) and (28) follow from the fact that, given that the cut-off type is defined by (26), the probability of a seller with cost x winning the contract is

$$\mu(x) = \begin{cases} 1, & x \leqslant x^* \\ [1 - F(x)]^{n-1}, & x^* < x \leqslant x_0 \\ 0, & x > x_0 \end{cases} \tag{29}$$

Equation (10) implies

$$p(x) = x + \frac{1}{\mu(x)} \int_x^{x_0} \mu(s)\, ds,$$

which yields (27).

With probability $(1-F(x^*))^{j-1}F(x^*)$, the buyer will contact j firms, while with probability $(1-F(x^*))^n$, he contacts all n firms. It follows that

$$\begin{aligned}
\phi = \sum_{j=1}^{n} &(p(x^*)+jc)(1-F(x^*))^{j-1}F(x^*) \\
&+ (1-F(x^*))^n nc + (1-F(x_0))^n z_0 \\
&+ \int_{x^*}^{x_0} p(y)\, n(1-F(y))^{n-1} f(y)\, dy \\
= p(x^*)&[1-(1-F(x^*))^n] \\
&+ \frac{c}{F(x^*)}[1-(n+1)(1-F(x^*))^n + n(1-F(x^*))^{n+1}] \\
&+ (1-F(x^*))^n nc + (1-F(x_0))^n z_0 \\
&+ \int_{x^*}^{x_0} yn(1-F(y))^{n-1} f(y)\, dy \\
&+ \int_{x^*}^{x_0} n \int_y^{x_0} (1-F(x))^{n-1}\, dx\, f(y)\, dy \\
= p(x^*)&[1-(1-F(x^*))^n] \\
&+ \frac{c}{F(x^*)}[1-(1-F(x^*))^n[n+1-n(1-F(x^*))-nF(x^*)]] \\
&+ z_0(1-F(x_0))^n + \int_{x^*}^{x_0} yn(1-F(y))^{n-1} f(y)\, dy \\
&+ nF(y) \int_y^{x_0} (1-F(x))^{n-1}\, dx \Big|_{x^*}^{x_0} + \int_{x^*}^{x_0} F(y)\, n(1-F(y))\, dy,
\end{aligned}$$

which implies (28) Q.E.D.

Just as the optimal direct incentive-compatible auction in the usual case of costless communication can be implemented as a sealed-bid or oral auction, so the optimal direct sequential incentive-compatible mechanism has its nondirect counterpart. The buyer in sequence invites potential sellers to submit price quotations. Bidder i, with production cost x_i, rationally bids

$p(x_i)$. The buyer either awards the contract to the first bidder who bids less than $p(x^*)$, or produces the item himself if no bid is less than $J(x_0)$. Note that, analogous with the usual equivalence between sealed-bid and oral auctions in the case of risk-neutral bidders, in the sequential auction the buyer is indifferent between receiving bids openly or in secret; in particular, the buyer gains nothing on average by informing a bidder about the best previous bid. The fact that the buyer rejects all bids and produces the good himself if the lowest bidder's production cost exceeds x_0 means that, as in the usual auction model, the buyer sets a reserve cost x_0 which is strictly less than his own production cost z_0 (since $z_0 = J(x_0) > x_0$), so that there is some probability of an inefficient outcome, with the buyer producing the item himself even though he has found a firm with a lower production cost.[10]

Implementing the optimal mechanism by a sequence of price quotations makes complete the search-theoretic interpretation. The cut-off production cost x^* defined by (26) implies a reservation-price rule. Let G represent the cumulative distribution of offered prices in the usual search formulation. Then with c as the unit search cost, the reservation price r is defined in the usual search model by[11]

$$c = \int_0^r (r - p)\, G'(p)\, dp = \int_0^r G(p)\, dp. \tag{30}$$

(This simply equates the marginal cost of taking one more observation with the marginal expected gain.) In the present model, the distribution of offered prices is $G(p) = F(J^{-1}(p))$, from Lemma 1.

THEOREM 10. *The reservation price r satisfies*

$$r = J(x^*). \tag{31}$$

[10] Since part (a) of Theorem 9 implies that the reserve price is the same for all potential sellers, the buyer is indifferent about whether or not the potential sellers know the order in which they are being approached. (Compare with the standard auction, in which the reserve price is independent of the number of bidders: Myerson [15], Riley and Samuelson [18].) However, the payment function varies. Equation (27) assumes the sellers do not know where they are in the order; if they did know this, payments would be the same on average, but the payment function would be different from (27).

[11] Note the distinction between the concepts of "reserve cost" (from auction theory) and "reservation price" (from search theory), which is about to be defined. On reserve prices, see McAfee and McMillan [11], Milgrom and Weber [12], Myerson [15], and Riley and Samuelson [18]. On reservation prices, see Lippman and McCall [8] and Rothschild [20].

Proof. Noting, by (26) and (30),

$$\int_0^r F(J^{-1}(p))\,dp = c = \frac{[F(x^*)]^2}{f(x^*)}, \tag{32}$$

we see that (31) is equivalent to

$$\int_0^{J(x)} F(J^{-1}(p))\,dp = \frac{[F(x)]^2}{f(x)}. \tag{33}$$

This holds since

$$\begin{aligned}
&\int_0^{J(x)} F(J^{-1}(p))\,dp - \frac{[F(x)]^2}{f(x)} \\
&= \int_0^x F(y)\,J'(y)\,dy - \frac{[F(x)]^2}{f(x)} \\
&= F(y)\,J(y)\Big|_0^x - \int_0^x J(y)\,f(y)\,dy - \frac{[F(x)]^2}{f(x)} \\
&= xF(x) - \int_0^x yf(y) + F(y)\,dy \\
&= xF(x) - yF(y)\Big|_0^x + \int_0^x F(y)\,dy - \int_0^x F(y)\,dy \\
&= 0.
\end{aligned} \tag{34}$$

Q.E.D.

The number of potential sellers, n, was assumed to be finite. However, taking limits, the buyer's expected total cost ϕ approaches the reservation price as the number of potential sellers becomes large. To see this, note from (28) that as $n \to \infty$,

$$\phi \to x^* + \frac{c}{F(x^*)} = x^* + \frac{F(x^*)}{f(x^*)} = J(x^*) = r. \tag{35}$$

Also, from (27), the expected payment received by the successful bidder, $p(y)$, approaches x^* as $n \to \infty$. This result was obtained by Riley and Zeckhauser [19].[12]

[12] In the context of the standard search model, a result analogous to this result (that, with perfect recall, the reservation price for search over a finite set of prices is the same as for search over an infinite set) was obtained by Landsberger and Peled [7] and Lippman and McCall [8].

Neither of the cut-off levels x_0 and x^* depends upon the amount of competition, n. Also, if $z_0 = \infty$, so that the buyer does not have the option of in-house production, $x_0 = \infty$, which implies that no bidder is ever immediately rejected: all bidders have a change of winning the final auction in the event of no bidder having a cost less than x^*. The assumed monotonicity of the function $J(x)$ implies that $[F(x)]^2/f(x)$ is monotonic increasing in x. Hence, from (26), the higher the search cost c, the higher the cut-off production cost x^*.

5. Conclusion

To summarize: whereas the monopsonist's optimal mechanism in the absence of communication costs is an oral or sealed-bid auction, when there are communication costs the optimal mechanism consists of reservation-price search followed, if the set of potential sellers is exhausted, by an auction. With an infinite set of potential sellers, the optimal mechanism is pure sequential search.[13]

The use of sequential mechanisms may, however, be undesirable when time matters and communication takes time. In this case, the buyer might wish to send signals to several potential sellers simultaneously, in order to reach a decision in less time. Our model generalizes to this case in a straightforward manner, with the caveat that a sequential mechanism may involve communications with several sellers simultaneously. Lemma 2 is consistent with this generalization. It may be possible to show, using a method analogous to our construction, that the search strategy of Morgan [13] and Morgan and Manning [14], involving samples of several observations, is optimal in the class of all mechanisms.

Appendix

We now give some more details of the extension of the Revelation Principle to the case of costly communication, and prove Lemma 2.

A principal-initiated, principal-centered mechanism allows an exchange of signals between the principal and some or all of the agents and, when

[13] In a complementary paper, Samuelson [21] analyzed an auction in which each bidder incurs a cost in submitting a bid. In the monopsonist's optimal auction, the buyer sets a reserve cost, which varies with the number of bidders. Such an extension is straightforward to implement in this model, since the effect is only to make the expected profit of the highest cost firm with a positive probability of winning equal to the cost of bidding. This means $\mu(x_m)$ will no longer be zero in Eq. (10).

this exchange of information has ended, culminates in a vector of decisions $(d_0, d_1, \ldots, d_n) \in D$.

Since communication is costly for the principal, there will be some agents with whom the principal will not want to communicate at any stage. It is notationally useful to introduce a nonsignal, ξ_ϕ; ξ_ϕ is interpreted as meaning no signal was sent.

The kth stage of a mechanism occurs in two parts. First, the principal sends signals $\xi_k = (\xi_k^i)_{i=1}^n$ to the n agents. (Many of these may be ξ_ϕ.) Denote the principal's earlier signals by ${}_{k-1}\xi = (\xi_j)_{j=1}^{k-1}$ and his earlier signals to agent i in particular by ${}_{k-1}\xi^i = (\xi_j^i)_{j=1}^{k-1}$ and his earlier signals to agent i in particular by ${}_{k-1}\xi^i = (\xi_j^i)_{j=1}^{k-1}$. The principal, having sent ξ_k to the agents, then receives from signals $\eta_k = (\eta_k^i, \ldots, \eta_k^n)$. Let ${}_{k-1}\eta = (\eta_j)_{j=1}^{k-1}$ and ${}_{k-1}\eta^i = (\eta_j^i)_{j=1}^{k-1}$. The signals ξ_k^i and η_k^i are members of some signal spaces given by the mechanism (note that ξ_ϕ is a member of both signal spaces).

When the principal chooses his kth-round signals ξ_k, he has previously chosen ${}_{k-1}\xi$ and observed ${}_{k-1}\eta$, and so he can condition his choice of ξ_k on these. To describe the choice of signal by the principal, let $\mu_k(\xi_k, {}_{k-1}\xi, {}_{k-1}\eta)$ be a probability distribution describing the choice of ξ_k.

Any agent's response η_k^i to the signal he receives, ξ_k^i, can be conditioned on ${}_k\xi^i$ and ${}_{k-1}\eta^i$ as well as the particular agent's type t_i. Let $\sigma_k^i(\eta_k^i, {}_k\xi^i, {}_{k-1}\eta^i, t_i)$ be the probability distribution governing the generation of the agent's reply η_k^i.

The fact that communication is principal-initiated means that σ_k^i must satisfy two requirements. First, a signal ξ_ϕ must be answered by ξ_ϕ, as ξ_ϕ means that no signal was sent by the principal to this agent. This implies that σ_k^i must satisfy:

$$\sigma_k^i(\eta_k^i, ({}_{k-1}\xi^i, \xi_\phi), {}_{k-1}\eta^i, t_i) = \begin{cases} 1, & \text{if } \eta_k^i = \xi_\phi \\ 0, & \text{otherwise.} \end{cases} \qquad \text{(A1)}$$

Second, let ${}_{k+1}\hat{\xi}^i = (\xi_i^i, \ldots, \xi_j^i, \xi_\phi, \xi_{j+1}^i, \ldots, \xi_k^i)$ and ${}_k\hat{\eta}^i = (\eta_1^i, \ldots, \eta_j^i, \xi_\phi, \eta_{j+1}^i, \ldots, \eta_{k-1}^i)$ for some $j \leq k$. Then

$$\sigma_{k+1}^i(\eta_0^i, {}_{k+1}\hat{\xi}^i, {}_k\hat{\eta}^i, t_i) = \sigma_k^i(\eta_0^i, {}_k\xi^i, {}_{k-1}\eta^i, t_i) \qquad \text{for any } \eta_0^i. \qquad \text{(A2)}$$

This is because ${}_{k+1}\hat{\xi}^i$ and ${}_k\hat{\eta}^i$ differ from ${}_k\xi^i$ and ${}_{k-1}\eta^i$ only by the addition of a nonsignal which, not being received, cannot affect the outcome of any future stage. Conditions (A1) and (A2) embody the fact that ξ_ϕ is merely a record-keeping device, and not a true signal.

Without loss of generality, the outcome ξ_ϕ from the distribution μ_k is used to denote the end of the communication process, where ξ_ϕ is the vector of nonsignals. Since it is possible that this point is reached before all of the agents have been communicated with, there must be some way of informing those agents who were not contacted that the process has ended. Introduce another fictitious signal ε, which the principal can, without cost, send to the agents he did not communicate with to tell them that the communication is finished.

Upon communication finishing at the kth stage, the principal chooses a decision d_0 according to the probability distribution $\mu_0(d_0, {}_k\xi, {}_k\eta)$, while each agent chooses his decision d_i given by the probability distribution $\sigma_0^i(d_i, {}_k\xi^i, {}_k\eta^i, t_i)$. The distribution σ_0^i must satisfy two conditions. First, since (as already noted) it must not be possible for the principal to use the mechanism so as to avoid incurring communication costs, there must be some exogenous decision $\hat{d}_i(t_i)$ which agent i takes if he has never been communicated with; that is,

$$\sigma_0^i(d_i, (\xi_\phi, ..., \xi_\phi, \varepsilon), (\xi_\phi, ..., \xi_\phi), t_i) = \begin{cases} 1, & \text{if } d_i = \hat{d}_i(t_i) \\ 0, & \text{otherwise.} \end{cases} \tag{A3}$$

Second, the addition of a nonsignal to the signals received by i must not change i's decision; that is, defining ${}_{k+1}\xi^i$ and ${}_k\eta^i$ as before,

$$\sigma_0^i(d_i, {}_{k+1}\hat{\xi}^i, {}_k\hat{\eta}^i, t_i) = \sigma_0^i(d_i, {}_k\xi^i, {}_{k-1}\eta^i, t_i), \qquad \text{for any } d_i. \tag{A4}$$

Equations (A3) and (A4) are analogous to (A1) and (A2). (Conditions (A1), (A2), (A3), and (A4) together can be taken to be a formal definition of principal-initiated communication.)

Let $\sigma^i = (\sigma_j^i)_{j=0}^\infty$; σ^i is a strategy for agent i. A message for agent i is $m_i = ({}_k\xi^i, {}_k\eta^i)$; denote M_i the set of messages. Formally, a mechanism consists of the signal spaces and the probability distribution $\mu = (\mu_j)_{j=0}^\infty$.

Any mechanism, when combined with a vector of strategies $\sigma = (\sigma^1, ..., \sigma^n)$, produces a distribution of outcomes $\zeta(d \mid t, \sigma)$. To solve for the distribution of outcomes, one generates the distribution of messages at the first stage, using μ_1 and σ_1. This is then used as an input into the second stage, using μ_2 and σ_2; and so on. This gives rise to a distribution of messages, which is then used in conjunction with μ_0 and σ_0 to produce a distribution of decision $d \in D$.

Agent i's expected utility is

$$V_i(\sigma) = \int_T \int_D U_i(t, d)\, \zeta(d \mid t, \sigma)\, P(t)\, dt\, dd. \tag{A5}$$

Consider a strategy vector σ. Let $\tilde{\sigma}^{-i} = (\sigma^1, ..., \sigma^{i-1}, \tilde{\sigma}^i, \sigma^{i+1}, ..., \sigma^n)$. Then σ is a Bayes–Nash equilibrium for the principal-initiated mechanism μ if and only if, for all $\tilde{\sigma}^i$ satisfying (A1), (A2), (A3), and (A4),

$$V_i(\sigma) \geqslant V_i(\tilde{\sigma}_i) \qquad \text{for all } i = 1, ..., n. \tag{A6}$$

The inability of an agent to communicate without prompting in a principal-initiated mechanism causes us to introduce yet another fictitious signal ξ_0. The principal, by sending ξ_0 to agent i, opens the communication and is the cue for the agent to respond with his type.

Formally, a mechanism is a *direct sequential* mechanism if the set of messages M_i contains only elements of the type (A7) or (A8). (Here it is supposed that the principal chooses to cease communication after the kth stage; and in (A7), the (ξ_0, t_i) element occurs at the jth component; that is, the jth stage of the sequential process):

$$({}_k\xi^i, {}_k\eta^i) = \begin{bmatrix} \xi_\phi & \xi_\phi \\ \vdots & \vdots \\ \xi_\phi & \xi_\phi \\ \xi_0 & t_i \\ \xi_\phi & \xi_\phi \\ \vdots & \vdots \\ d_i & \xi_\phi \end{bmatrix} \equiv \gamma_j; \tag{A7}$$

$$({}_k\xi^i, {}_k\eta^i) = \begin{bmatrix} \xi_\phi & \xi_\phi \\ \vdots & \vdots \\ \xi_\phi & \xi_\phi \\ \varepsilon & \xi_\phi \end{bmatrix}. \tag{A8}$$

With the message (A7), the principal communicates with agent i at the jth stage, asking and being told i's type. The principal then, at the kth and last stage, signals to i that the communication is over by suggesting the decision d_i for agent i. With the message (A8), the principal never communicates with i except at the end to inform the agent, without incurring any cost of communication, that the process is ended.

An agent is honest if

$$\sigma_j^i(t_i, \xi_0, t_i) = 1 \qquad \text{for all } t_i. \tag{A9}$$

Equation (A9) implies, using (A2), that agent i correctly reports his type when asked (that is, upon receiving the message ξ_0 as in (A7)). An agent is obedient if, with γ_j given by (A7),

$$\sigma_0^i(d_i, \gamma_j, t_i) = 1 \qquad \text{for all } t_i; \tag{A10}$$

that is, the agent takes the decision recommended by the principal.

Let $\omega \in \Omega$ represent an ordered set of agents communicated with. On occasion, in a loose but not misleading use of notation, we shall refer to agent $i \in \omega$, meaning that $\omega = (i_1, ..., i_{j-1}, i, i_{j+1}, ..., i_k)$.

Proof of Lemma 2. The proof begins by constructing μ^* from μ, σ. It is then demonstrated that μ^* is incentive compatible. It will be clear from the construction that, for any t, the distribution of decisions and agents communicated with is unchanged. The construction is by induction. Denote by v_k the set of agents communicated with in the direct sequential mechanism μ^* by the end of the kth stage.

The base of the induction requires

$$\mu_1^*(\xi_\phi, ..., \xi_\phi, \xi_0, \xi_\phi, ..., \xi_\phi) = \sum_{\xi_1^i \neq \xi_\phi} \mu_i(\xi_\phi, ..., \xi_\phi, \xi_1^i, \xi_\phi, ..., \xi_\phi). \quad \text{(A11)}$$

Let v_0 be the empty tuple. This provides the base of the induction.

Now suppose $({}_{k-1}\xi, {}_{k-1}\eta)$, an internal vector to μ^*, has been given. The following steps represent the internal workings of the mechanism.

1. Choose ξ_k from the distribution $\mu_k(\xi_k, {}_{k-1}\xi, {}_{k-1}\eta)$. If $\xi_k = \xi_\phi$, go to 4. If there is no $i_k \in \omega$ with ${}_k\xi^{i_k} = ({}_{k-1}\xi_\phi, \xi_k^{i_k})$ and $\xi_k^{i_k} \neq \xi_\phi$, go to 3. Otherwise proceed to 2.

2. To be here, there must be an agent i_k receiving his first signal. Send i_k the signal ξ_0. Let $v_k = (v_{k-1}, i_k)$. i_k responds with this type t_{i_k}. Go to 3.

3. Since t_i is known for all types who have been sent any message other than ξ_ϕ, we may operate σ_k^i on $({}_k\xi^i, {}_{k-1}\eta^i)$ and generate a signal, internal to the mechanism, η_k^i, which is added to ${}_{k-1}\eta^i$ to produce ${}_k\eta^i$. For other $i \notin v_k$, ${}_k\xi^i = \xi_\phi$, so set $\eta_k^i = \xi_\phi$. Return to 1.

4. The last value of ξ_k was ${}_k\xi_\phi$. For all $i \in v_k$, draw a decision d_i from the distribution $\sigma_0^i(d_i, {}_k\xi^i, {}_k\eta^i, t_i)$. Send this decision d_i to i, and send ε to all other agents. Draw a decision d_0 from the distribution $\mu_0(d_0, {}_k\xi, {}_k\eta)$.

Note that the resulting mechanism μ^* is direct and sequential. By construction, $v_k = \omega$. Since the same distributions are used, the same distribution of outcomes occurs for each $t \in T$, provided the agents are honest and obedient. Finally, to see that this mechanism is incentive compatible, suppose that at state j agent i responds with type t_i' when his true type is t_i. Then, in the original mechanism, the σ_j^i arising from type t_i' must dominate the σ_j^i arising from the type t_i. Thus, σ^i was suboptimal, contrary to hypothesis. Similarly, if the agent does not choose d_i as suggested by the principal, there must be a preferred decision d_i', contrary to the hypothesis that σ_0^i was optimal. Q.E.D.

If it is desirable to include real time and discounting in the model, the definition of sequential mechanisms must be modified to allow asking several agents their type simultaneously. The only modification necessary to the proof is allowing the principal to let real time pass in the computation of ε_k, in step 3, so that μ^* involves the same timing as μ.

Note that this proof is essentially the same as the proof of the usual Revelation Principle (Myerson [16]), except that the fact that communication is costly means that care must be taken over the timing of messages: in particular, the proof must keep track of exactly which agents have been communicated with, and which have not, at any stage. This is why the ficitious signals ξ_ϕ, ξ_0, and ε are needed.

References

1. K. Burdett and K. L. Judd, Equilibrium price disperson, *Econometrica* **51** (1983), 955–969.
2. J. A. Carlson and R. P. McAfee, Discrete equilibrium price dispersion, *J. Polit. Econ.* **91** (1983), 480–493.
3. P. A. Diamond, A model of price adjustment, *J. Econ. Theory* **3** (1971), 156–168.
4. J. R. Fox, "Arming America: How the U.S. Buys Weapons," Harvard Univ. Press, Cambridge, MA, 1974.
5. E. J. Green, On the difficulty of eliciting summary information, *J. Econ. Theory* **32** (1984), 228–245.
6. J. Green and J.-J. Laffont, "Limited Communication and Incentive Compatibility," Discussion Paper No. 919, Harvard University, 1982.
7. M. Landsberger and D. Peled, Duration of offers, price structure, and the gain from search, *J. Econ. Theory* **16** (1977), 17–37.
8. S. A. Lippman and J. J. McCall, The economics of job search: A survey, *Econ. Inquiry* **14** (1976), 155–189.
9. E. Maskin and J. Riley, Monopoly with incomplete information, *Rand J. Econ.* **15** (1984), 171–196.
10. R. P. McAfee and J. McMillan, Bidding for contracts: A principal-agent analysis, *Rand J. Econ.* **17** (1986), 326–338.
11. R. P. McAfee and J. McMillan, Auctions and bidding, *J. Econ. Lit* **25** (1987), in press.
12. P. R. Milgrom and R. J. Weber, A theory of auctions and competitive bidding, *Econometrica* **50** (1982), 1089–1122.
13. P. B. Morgan, Search and optimal sample sizes, *Rev. Econ. Stud.* **50** (1983), 659–676.
14. P. B. Morgan and R. Manning, Optimal search, *Econometrica* **53** (1985), 923–944.
15. R. B. Myerson, Optimal auction design, *Math. Operations Res.* **6** (1981), 58–73.
16. R. B. Myerson, Optimal coordination mechanisms in generalized principal-agent problems, *J. Math. Econ.* **10** (1982), 67–81.
17. J. F. Reinganum, Strategic search theory, *Int. Econ. Rev.* **23** (1982), 1–18.
18. J. Riley and W. Samuelson, Optimal auctions, *Amer. Econ. Rev.* **71** (1981), 381–392.
19. J. Riley and R. Zeckhauser, Optimal selling strategies: When to haggle, when to hold firm, *Quart. J. Econ.* **98** (1983), 267–290.
20. M. Rothschild, Models of market organization with imperfect information: A survey, *J. Polit. Econ.* **81** (1973), 1283–1308.
21. W. F. Samuelson, Competitive binding with entry costs, *Econ. Letters* **17** (1985), 53–58.
22. L. L. Wilde, Equilibrium search models as simultaneous move games, mimeo, California Institute of Technology, 1985.

Printed by Catherine Press, Ltd., Tempelhof 41, B-8000 Brugge, Belgium