A General Solution To King Solomon’s Dilemma*

Motty Perry
Department of Economics and
The Center for Rationality and
Interactive Decision Theory
Hebrew University of Jerusalem
Jerusalem 91904, Israel

Philip J. Reny
Department of Economics
University of Pittsburgh
Pittsburgh, PA 16046

September 28, 1994

(Revised July 25, 1996)

*We have benefited from helpful conversations with Murali Agastya, George Mailath, Bob Rosenthal, Dov Samet, Larry Samuelson, and Shmuel Zamir. The second author thanks the Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem for its generous hospitality during a visit on which this research was conducted. Both authors gratefully acknowledge support from the U.S.-Israel BSF (grant #95-00023/1).

1. Introduction

A classical situation involving disagreement between two parties is "King Solomon's Dilemma". The nature of Solomon's dilemma begins with a dispute between two women in which each claims to be the mother of a certain child. Of course, Solomon wishes to give the child to the rightful mother *at no cost to her*. The difficulty is that although Solomon knows that one of them is the mother, he does not know which one.[1]

In a convincing application of the theory of implementation, Glazer and Ma (1989) provide a strikingly simple and elegant solution to Solomon's dilemma for the case in which it is commonly known among all three parties that the rightful mother values possession of the child at a dollars, and the impostor values possession of the child at b dollars, where $a > b$. Their paper also contains examples of King Solomon-like dilemmas relevant to economics.

Moore (1991) modified the Glazer-Ma mechanism to nicely accommodate the more general case in which it is commonly known that only the women know the values of a and b, while Solomon knows only that the true mother values possession of the child strictly more than the impostor does.[2]

Our objective here is to complete the task of resolving Solomon's dilemma by removing the remaining informational restriction while maintaining simplicity in the implementing mechanism. We shall assume only that it is commonly known both that the women know who the rightful mother is and that she values possession of the child strictly more than does the impostor. It is shown that a second-price sealed-bid all-pay auction with the winner having an ex-post option to quit solves King Solomon's dilemma in iteratively undominated strategies.

The full force of iterative dominance is not needed here. Only four rounds of elimination are required. This is of some significance since earlier rounds of elimination can be justified more easily than later ones. Later rounds can be justified only if the players know that the previous eliminations have been made. Consequently, fewer rounds of elimination correspond to less stringent assumptions about the players' mutual knowledge and therefore render the solution more compelling.

The remaining two sections present the model and solution. For ease of exposition, the formal details are kept to a minimum, although the analysis can easily be carried out within a standard model of knowledge such as in Aumann (1976).

2. The Model

---

[1] Without the "zero cost" constraint, a standard Vickrey auction would suffice in allocating the child to the true mother since she values the child more than does the impostor.

[2] In an appendix, Glazer and Ma (1989) actually provide another mechanism to cover this case as well. However, Moore's (1991) mechanism is much simpler.

There are two agents, *A* and *B*. A single object is to be allocated at no cost to the agent who values it most. We maintain the following assumptions, each of which is common knowledge between *A* and *B*:

(i) the agents' values are non negative and distinct,
(ii) each agent knows his own value and each agent knows which of them has the higher value,
(iii) neither agent rules out the true value of the other agent,
(iv) the low value agent places a finite upper bound on the other agent's value, and
(v) each agent's payoff of obtaining the object at price *p* when its value to that agent is *v*, is *v*-*p*, while the payoff associated with paying *p* and not receiving the object is -*p*.

Requiring the agents' values to be non negative is not essential. For example, problems in which both agents' values are non positive, such as the allocation of a dump site to one of two municipalities, can be accommodated by defining the object as "the right to allocate" the dump site.

Assumption (iii) might strike the reader as being very strong. In our view this is not the case. For to violate (iii) an agent must rule out the truth which is tantamount to drawing a definite conclusion when no such conclusion can possibly be (definitively) drawn. In any event, the present informational assumptions are substantially weaker than those in Glazer and Ma (1989) and Moore (1991). The presence of a known upper bound on the value of the object embodied in (iv) can be dispensed with at the cost of complicating the mechanism slightly.

The Mechanism: The agents will participate in a second-price sealed-bid all-pay auction *with an option*. The option is that after the bids are revealed to the agents, the winner (highest bidder) can either choose to stick with his bid (in which case he receives the object and both bidders pay the second-highest bid), or he can choose to *quit* and give the object to the other agent in which case no payments are made by either agent. If the two bids are identical, then the object is sold to one of them (determined by the toss of a fair coin) at a price equal to the common bid. In this case the other agent pays nothing.

Although we shall consider only the case of two agents, the analysis easily extends to the *n*-agent case. When there are $n > 2$ agents, the mechanism works as above except that the option to quit becomes a sequential one. For example, if there are three agents and the bids are 10, 20 and 30, then the 30 bidder has the option to buy at 20 (with all agents paying 20) or to quit. If he quits, the 20 bidder then has the option to buy at 10 (with all bidders other than those who have previously quit paying 10) or to quit. If he quits, the 10 bidder receives the object for free and no one pays anything. In addition, for the *n*-agent case, the second part of item (ii) is replaced by : "the highest-value agent knows that his value is strictly higher than all other agents' values." This implies in particular that the highest-value agent knows who he is, and that he may be the only one who knows who the highest-value agent is.

3. The Solution

We now show that the mechanism above implements the desired outcome in iteratively undominated strategies.

A *strategy* is a function mapping an agent's information (i.e. his own value, his knowledge of the other agent's value, etc.) into a non negative bid and a decision function. The decision function provides for each pair of bids in which the agent's bid is winning, a decision to either quit or not.

A strategy, s, for A *dominates* another of A's strategies, s', against a subset, T, of B's strategies, if for every t in T, s yields at least as high a payoff for A as s' against t regardless of the two agents' information, and a strictly higher payoff against some t in T for some information the agents' may possess. A dominated strategy for agent B is defined similarly.

The set of *iteratively undominated strategies* are those that remain after eliminating in successive rounds all strategies that are dominated against those present in that round. See, for example, Moulin (1979, 1981).

In the analysis below, we eliminate particular dominated strategies in each round, not checking whether or not we have eliminated *all* dominated strategies in each round. However, it is not difficult to show that the unique outcome that we isolate can also be obtained through eliminations which are "nice" in the sense of Marx and Swinkels (1994). For the details, consult Perry and Reny (1996). Consequently, the order of elimination is irrelevant. In particular, the same outcome would result if *all* dominated strategies were eliminated in every round.

Without loss of generality, we assume that agent A values the object more than B. Throughout the analysis below, p denotes a bid by A and q denotes a bid by B. We now begin the rounds of elimination.

Round 1: For each agent, eliminate every strategy such that given the agent's information, which includes his value, and the bid specified by the strategy, the strategy also specifies quitting (resp., buying the object) if the agent's bid is winning and the second highest bid is below (resp., above) his value.

Round 2: Eliminate all strategies for A in which he bids above his value, a. All such bids are dominated by bidding his value. To demonstrate this, we consider below all possible cases. It is useful to recall that by assumption

(*) A knows that his own value, a, strictly exceeds B's value.

(a) $q > p > a$: B wins the auction whether A bids p or a, and in both cases B exercises the option to quit. Hence, by (*), A knows this and so is indifferent between bidding p and a.

(b) $q = p > a$: By bidding p, agent A, with probability one-half, obtains the object for a price of $p > a$. However, a bid equal to a would render B the winner. B would then take the option to quit giving A the object for free. By (*), A knows this so that A strictly prefers the bid a over p.

(c) $p > q > a$: If A bids p, then A wins the auction and takes the option to quit. But if A bids a, then B wins the auction and takes the option to quit. By (*) A knows this so that bidding a is strictly better for A than is bidding p.

(d) $p > a = q$: If A bids p, then A wins the auction and obtains a payoff of zero whether or not he chooses to quit. If A bids a, then with probability one-half he is the winner and again receives a payoff of zero, and with probability one-half B is the winner in which case A neither pays any money nor receives the object. Hence, A is indifferent between bidding a and p.

(e) $p > a > q$: A wins the auction whether he bids a or p. Hence A is indifferent between bidding a and p.

Since there are strategies for B remaining after round 1 which can lead to any of the cases (a)-(e) above, we conclude that it is dominant for A to submit a bid less than or equal to his value.

Round 3: Eliminate all remaining strategies for B except those in which he chooses a bid that he knows is above A's value. That B places some upper bound on A's value is guaranteed by assumption, and this is common knowledge. Note that if agent B makes such a bid, then because he does not rule out A's true value, his bid is indeed above A's true value.

To see that these strategies are dominated, consider first all remaining possibilities at this stage in the elimination process. Recall that after round 2, agent A does not bid above his value, so that p £ a. The remaining possibilities then are:

(a) $q < p$ £ a: If B bids q, then A wins the auction and chooses to buy the object at price q. B must then also pay q obtaining a non positive payoff. If instead B bids above A's value, then B wins the auction and so is guaranteed a non negative payoff since he can exercise the option to quit. Moreover, for values of p below B's value of the object, B strictly prefers to bid above A's value since he will obtain a strictly positive payoff by purchasing the object after winning the auction.

(b) $q = p < a$: If B bids q, then with probability one-half B must buy the object at price p. Bidding above A's value guarantees B the option of buying the object at price p. The latter is strictly better for B whenever p differs from B's value and equally good otherwise.

(c) $q = p = a$: If B bids q, then with probability one-half B must buy the object at price $p = a$ which is above his value, while bidding above A's value guarantees B a non negative payoff.

(d) $q > p$, p £ a: B wins the auction whether he bids above A's value or bids q. Hence B is indifferent between the two.

To conclude that B submits a bid that he knows is above A's value, we must ensure that if he did not, then at least one of cases (a)-(c) can actually occur (since these are the cases in which the proposed dominating strategy is strictly better). But this is straightforward, since if he does not submit a bid that he knows is above A's value, then any one of cases (a)-(c) can occur, since agent A's value, a, might then be greater than or equal to B's bid, q.

Thus, agent *B* submits a bid that he knows is above *A*'s value. Because agent *B* does not rule out the truth, his bid is then, in fact, above *A*'s value.

Round 4: Eliminate all remaining strategies for *A* except those in which he chooses a bid that he knows is above *B*'s value. Note that agent *A* can accomplish this by bidding his own value since he knows that his value is highest. The reason that these strategies are dominated follows.

From round 3, we have that *B*'s bid exceeds *A*'s value. Since at this stage *A*'s bid does not exceed his own value, *A* knows that *B* will win the auction. Consequently, by bidding above *B*'s value, *A* guarantees himself the object for free, since *A*'s bid is then certain to induce *B* to quit after *B* wins the auction. On the other hand, among those bids remaining (i.e. those which do not exceed *A*'s value), were agent *A* to choose one that might not be above *B*'s value, *A* runs the risk that *B* would then choose to purchase the object. Agent *A* then would not obtain the object and would also pay the amount of his bid.

We conclude that *A* submits a bid that he knows is above *B*'s value. Because agent *A* does not rule out the truth, his bid is then, in fact, above *B*'s value.

So at this point in the elimination process all remaining strategies are such that *A*'s bid is not above his own value and it is above *B*'s value, while *B*'s bid is above *A*'s value. Consequently, all remaining strategies yield the same outcome, namely that *B* wins the auction and chooses to exercise the option to quit. No further eliminations are possible. Thus *A* receives the object and neither agent makes any payment.

References

Aumann, R.J. (1976): " Agreeing to Disagree," *The Annals of Statistics*, 4, 1236-1239.

Glazer, J., and C-T. Ma (1989): "Efficient Allocation of a "Prize"--King Solomon's Dilemma," *Games and Economic Behavior* 1, 222-233.

Marks, L. and J. Swinkels (1994): "Order Independence & Iterated Weak Dominance," mimeo, Department of Economics, Northwestern University, and forthcoming in *Games and Economic Behavior*.

Moore, J. (1991): "Implementation in Environments with Complete Information," STICERD working paper TE/91/235, London School of Economics and Political Science.

Moulin, H. (1979): "Dominance-Solvable Voting Schemes," *Econometrica*, 47, 1337-1351.

------------ (1981): *Game Theory for the Social Sciences*. New York: New York University Press.

Perry, M. and P. J. Reny (1996): "A General Solution to King Solomon's Dilemma," mimeo, Department of Economics, University of Pittsburgh.

Appendix I

In this appendix we provide a formal model of knowledge along the lines of Aumann (1976).

Let $\boldsymbol{W}$ denote the set of states of the world and let $\boldsymbol{P}_A$ and $\boldsymbol{P}_B$ denote agent A's and B's information partitions of $\boldsymbol{W}$ respectively. Let $\boldsymbol{P}_i(\boldsymbol{w})$ denote the element of $\boldsymbol{P}_i$ containing those states that i does not distinguish between (or rule out) when the true state is ω. Also, let $\Psi:\Omega\rightarrow\mathbf{R}^2_{++}$ be a mapping taking states of the world into a value of the object for each agent. We maintain the following assumptions for all $\omega, \omega'\in\Omega$ and $i = A, B$:

1. $\omega\in\Pi_i(\omega)$
2. If $\Psi(\omega)= (x_A,x_B)$ then
   (a) $x_A {}^{1} x_B$
   (b) $\omega'\in\Pi_i(\omega) \Rightarrow \Psi_i(\omega')=x_i$
   (c) $x_A > x_B$, $\omega'\in\Pi_i(\omega)$ and $\Psi(\omega')=(y_A,y_B) \Rightarrow y_A > y_B$.
3. $\Psi(\Pi_i(\omega))$ is bounded.

Assumptions 1-3 express, respectively that: 1. neither agent rules out the truth; 2(a) the agents' values are distinct, (b) each agent knows his own value, (c) each agent knows whose value is larger; and 3. each agent places a finite upper bound on the other's value at each state of the world.

A *strategy* for an agent is a function from states of the world into a non negative bid and a decision function. The decision function specifies for each pair of bids in which the agent's bid is winning, a decision to either quit or not. The agent's strategy must be measurable with respect to his information partition.

Given strategies s and t for agents A and B respectively, let $u(s(\omega), t(\omega)|\omega)$ denote A's payoff from the auction when the state is $\omega\in\Omega$. Let s' be another strategy for A and let T be a subset of strategies for B. Then we say that *s dominates s*' for A against T, if $u(s(\omega),t(\omega)|\omega) \geq u(s'(\omega), t(\omega)|\omega)$ for all $\omega\in\Omega$ and all $t\in T$, with at least one such pair ω and t yielding a strict inequality. Dominance is similarly defined for agent B.

With these definitions, the steps taken in the main text can be applied equally well here and the result is the same. Our mechanism implements the desired outcome in iteratively undominated strategies whenever the situation can be modeled as above. We now give two such examples. The examples are also meant to illustrate the permissiveness of our informational assumptions.

Example 1 (Glazer and Ma (1989)): Fix $a > b$. Let $\Omega=\{(a,b), (b,a)\}$, $\Pi_A=\Pi_B=\{\{(a,b)\},\{(b,a)\}\}$ and $\Psi(\omega)=\omega$ for all $\omega\in\Omega$.

Example 2 (Moore (1991)): Let $\Omega=\{(a,b)\in\mathbf{R}^2_{++}| a \neq b\}$, $\Pi_A=\Pi_B= \{\{(a,b)\}|(a,b)\in\Omega\}$ and $\Psi(\omega)=\omega$ for all $\omega\in\Omega$.

Appendix II

Not all eliminations in the text are “nice” in the sense of Marks and Swinkels (1994). In this appendix, it is shown that the single outcome that we derived in the text in four rounds of elimination can be derived in five rounds by employing only “nuice” eliminations. Consequently, by Marks and swinkels (1994), the order of elimination does not matter.

There are two eliminations that are not “nice” in the text. They occur in rounds 2 and 3. We now show how these can be taken care of.

Consider first the elimination carried out in round 2. In case (d) there, agent *B* is not indifferent between *A*’s bid of *p* and *A*’s bid of *a*, while agent *A* is. Therefore, this elimination is not, formally, “nice”. However, the elimination can be rendered “nice” by dominating *A*’s bid of *p* not by the bid of *a*, but rather by the mixed bid in which *A* chooses a bid uniformly from the nondegenerate interval (*a*,*p*). Of course, bidding *a* is much simpler and more natural than employing the suggested mixed bid, and this is why it appears in the elimination process. The mixed bid should be thought of as a technical device used to demonstrate that the eliminations up to this point are “nice”.

Consider now round 3. In case (b) there, if $q = b$, then agent *B* is indifferent between bidding *q* and bidding above *A*’s value, but agent *A*’s payoff is affected by *B*’s choice. Hence, this elimination is not “nice”. However, this can be taken care of by not eliminating in round 3 the strategies in which agent *B* bids his value. The strategies eliminated in round 4 will nonetheless remain the same and adding a fifth round of elimination will get rid of the strategies in which agent *B* bids his value. All of these eliminations *are* “nice” and the end result is a single outcome (and the one we isolate in the text in four rounds). Consequently, by Marx and Swinkels (1994) the order of elimination does not matter.