On the Value of a Well Chosen Bit to the Seller in an Auction

Bruce Hajek
Department of Electrical and Computer Engineering
and the Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Email: b-hajek@uiuc.edu

***Abstract*— A central fact in the theory of optimal auction design is that the seller of a single object in an auction with n bidders, having independent, random valuations, typically cannot extract the full maximum value of the object from the buyers. We show that if the seller has access to a single bit of information, even if noisy, then the seller can extract full value. The work is meant to explore the use of information measures in mechanism design problems.**

I. Introduction

There is an extensive theory of optimal auctions and design of allocation mechanisms, including the seminal papers of Myerson [6] and Riley and Samuelson [8], and recent books [2–4]. In one version of the problem, a single indivisible object is to be sold. Each of n bidders or potential buyers submits a bid or signal to the seller, and then one of the bidders is selected to buy the object and make a payment for it. Examples include the sale of a rare painting, or the sale of a block of communication spectrum. Elements of game theory enter, because the bidders are assumed to act strategically, in trying to maximize their own payoffs. Many theories of bounded rationality have been proposed over the past decade or two in order to account for the fact that the agents involved (bidders and seller) may have limited computation power. In some situations, there may also be limitations on communication. It may therefore be interesting to assess the tradeoffs between optimality or efficiency of allocation mechanisms, and the information requirements. As a step in this direction, this paper explores the value of a possibly noisy single bit of information to a seller in the context of optimal auction design.

II. Overview of Optimal Auctions

Consider the auction of a single indivisible object by a single seller to n bidders. Suppose the state of bidder i is $s_i \in \{1, \cdots, m_i\} = \mathcal{M}_i$, with probability mass function $\pi_i(s_i) > 0$ for $s_i \in \mathcal{M}_i$. Also, suppose each bidder i has a nonnegative valuation function w_i defined on $\mathcal{M}_i$, so that $w_i(s_i)$ is the value of the object to bidder i. Suppose the states $s_1, \cdots, s_n$ are mutually independent. Each bidder and the seller know the probability mass functions π_i and valuation functions of all the bidders, but the state s_i is private information of bidder i. Assume the object has value $w_0 \equiv 0$ to the seller. Each bidder i reports a state s'_i. A *mechanism* (p, x) is given by an allocation probability distribution $p = (p_i(s') : 1 \leq i \leq n)$ and average charges $(x_i(s') : 1 \leq i \leq n)$ to the bidders, for each vector of reported states $s' = (s'_1, \cdots, s'_n) \in \mathcal{M} = \mathcal{M}_1 \times \cdots \times \mathcal{M}_n$. If $s \in \mathcal{M}$ and $1 \leq i \leq n$, then s_{-i} denotes the vector $(s_j : j \neq i)$. A mechanism is *Bayes-Nash incentive compatible* (BN-IC) if for each i, if all other bidders report their true values (i.e. $s'_{-i} = s_{-i}$) then bidder i maximizes her own expected return by also telling the truth. A BN-IC mechanism is said to be *interim individually rational* (IIR) if, under truth telling by all bidders,

$$E[\text{net gain of bidder } i|s_i] \geq 0.$$

The maximum possible gain that can be created by allocation of the object is $E[\max_i\{w_i(s_i)\}]$. This gain is created if and only if the allocation is efficient, meaning the object is sold to a bidder with maximum valuation. A mechanism is called *full seller extraction* (FSE) if the seller's expected payoff is $E[\max_i\{w_i(s_i)\}]$.

One could easily imagine auction mechanisms in which the signal sent from a buyer i takes values in a space different from $\mathcal{M}_i$, and the buyers are allowed to use mixed strategies–randomizing their bids. However, the *revelation principle* of auction theory [6] implies that, without loss of optimality, the seller can use auctions in which the Bayes-Nash equilibrium which maximizes the expected payment to the seller, is such that the signal from buyer i is a truthful report of the buyer's state in $\mathcal{M}_i$. That is, there is no loss of optimality in seeking BN-IC mechanisms, based on revelation of state.

With this formulation, the search for a mechanism which maximizes the revenue to the seller can be expressed as a linear programming problem in the x and p variables. The expected revenue and the BN-IC and IIR constraints are linear in these variables. Typically there are no FSE, BN-IC, IIR allocation mechanisms, or equivalently, no FSE IIR allocation mechanisms.

A simple example of an auction is the case of two identical buyers, each valuing the object at H with probability p, and L with probability $1-p$, where $0 \leq L < H$, and $0 < p < 1$. This corresponds to $n = 2$ and, for $i \in \{1, 2\}$, $m_i = 2$, $\pi_i = (1-p, p)$, and $w_i = (L, H)$. The problem is simple enough that the linear equations for optimal auction design (see [5, 6]) can be solved without the aid of a computer. For some parameter values the seller can extract more revenue on the average if she is permitted to not sell the object to either bidder. In case the seller must sell the object, the optimal mechanism is to sell the object at the price $b_H = \frac{H+(1-p)L}{2-p}$ if at least one bidder reports the high state (i.e. state 2, for value H) and at price L if both bidders report the low state. The object is sold to the bidder with the higher bid if the bids are unequal, and to either bidder with probability one half if the bids are equal. The price b_H is strictly smaller than H. If it weren't, a high state bidder would have incentive to report the low state to the seller. In fact, the price b_H is precisely the value such that

$$(\frac{p}{2} + 1 - p)(H - b_H) = (\frac{1-p}{2})(H - L),$$

which means that a high state bidder would receive the same expected payoff for reporting the high state to the seller as reporting the low state to the seller. In the notation of this paper, the optimal mechanism is given by

$$\begin{array}{ll} p(2,2) = (\frac{1}{2}, \frac{1}{2}) & x(2,2) = (\frac{b_H}{2}, \frac{b_H}{2}) \\ p(2,1) = (1, 0) & x(2,1) = (b_H, 0) \\ p(1,2) = (0, 1) & x(1,2) = (0, b_H) \\ p(1,1) = (\frac{1}{2}, \frac{1}{2}) & x(1,1) = (\frac{L}{2}, \frac{L}{2}) \end{array}$$

Another way to view the optimal mechanism is that the object is sold to the bidder offering the higher price, at the price bid (i.e. a first price auction), with ties broken by a fair coin flip, but only bid values L and b_H are accepted. The mechanism is not FSE because the buyers retain excess value $H - b_H$ whenever at least one buyer has the high state. The amount of expected value that the seller does not extract is thus $(2p - p^2)(H - b_H)$, which simplifies to $p(1-p)(H-L)$. In a sense, $p(1-p)(H-L)$ is the expected cost to the seller, sometimes called *information rent*, for not knowing the private values of the buyers.

If the seller has the option of withholding the object, the above mechanism is still optimal if $L \geq pH$. If instead, $L < pH$, the optimal mechanism if withholding is permitted is for the seller to sell at price H if at least one bidder reports the high state, to a seller reporting the high state. If both bidders report the low state, then the seller withholds the object and receives no payment. If $L > 0$, the mechanism is not FSE because the seller fails to extract the value L when both buyers are low state. The expected amount of value not extracted by the seller is $L(1-p)^2$.

Another indication that FSE is typically not possible is offered by the case of auctions with continuously distributed valuations. Suppose f is a probability density function on a finite interval $[a, b]$ and F is the corresponding cumulative distribution function, such that $t - \frac{1-F(t)}{f(t)}$ is strictly increasing. For example, f could be the uniform density for $[a, b]$. If the value of the object to a buyer i has the density function f, then the optimal auction is Vickrey's second price auction [6]. The object is sold to the highest bidder with the price equal to the second highest bid. The dominant strategy for any buyer in a second price auction is for her to bid her true state [9], so that the mechanism is BN-IC. The mechanism is not FSE, because the seller receives only the second highest valuation. Auctions with nonsymmetric buyers having continuous valuations are also typically not FSE [6].

III. On the Value of a Reliable Bit

So far we assume the seller has no private information. Consider next a situation in which the seller has access to a binary valued piece of private side information b, where b is a random variable with values in $\{0, 1\}$. For example, one could think that a helpful genie or spy could report a binary valued function of $s = (s_1, \ldots, s_n)$ to the seller. While b is allowed to be correlated with s, to avoid trivialities b is not otherwise allowed to depend on the vector of reported values s'. That is b and s' are conditionally independent given s. For example, it is not allowed to take $b = 1$ if and only if the bidder with the largest bid lies. Given the individual distributions π_i, the joint distribution of s and b is completely determined/specified by the function $\phi : \mathcal{M} \times \{0, 1\} \to [0, 1]$ representing the conditional probability $P[b|s] = \phi(s, b)$. Note that $\phi(s, 0) + \phi(s, 1) = 1$ for all s.

Assume that the seller and the bidders know that b is available to the seller, and that they know the joint distribution of s and b. The question we address is how much value can the seller extract due to the existence and use of b, for a best choice of ϕ.

Proposition 3.1: Suppose $\log_2 m_i \leq \prod_{j \neq i} m_j$ for each i. Then there is a choice of ϕ such that there exists an FSE, IIR, BN-IC mechanism for a seller with binary private information b.

The proof of the proposition is based on using a function $u : \mathcal{M} \to \{0,1\}$ with a certain property as a hashing function. The function u is said to be sensitive to its ith coordinate if for any two distinct values s_i and s_i', there is a value of s_{-i} such that $u(s_i, s_{-i}) \neq u(s_i', s_{-i})$.

Lemma 3.2: Under the condition of Proposition 3.1, there is a function $u : \mathcal{M} \to \{0,1\}$ which is sensitive to every coordinate.

Proof Without loss of generality, suppose $m_1 \geq m_2 \geq \cdots \geq m_n$. By assumption, $m_1 \leq 2^{m_2 \cdots m_n}$. Assume in addition that $m_1 = m_2$. The lemma will first be proved under this assumption, and then proved in general. Define u (using the notation $a \wedge b = \min\{a, b\}$) for $s \in \mathcal{M}$ by

$$u(s) = \begin{cases} 1 & \text{if } s = (a, a, a \wedge m_3, a \wedge m_4, \cdots, a \wedge m_n) \\ & \quad\text{for some } a \text{ with } 1 \leq a \leq m_2 \\ 0 & \text{else} \end{cases}$$

We show that u is sensitive to every coordinate. To that end, fix i with $1 \leq i \leq n$ and let s_i and s_i' be distinct values in $\mathcal{M}_i$. Let $s_{-i} = (s_j : j \neq i)$ be determined by $s_j = s_i \wedge m_j$ for $j \neq i$. Then $u(s_i, s_{-i}) = 1$ and $u(s_i', s_{-i}) = 0$. Therefore u is sensitive to coordinate i, and hence to any coordinate. The lemma is thus true under the assumption that $m_1 = m_2$. (The above construction wouldn't work if $m_1 > m_2$, because if $i = 1$ and $s_1 \geq m_2$ and $s_1' \geq m_2$ then we would have $u(s_i, s_{-i}) = u(s_i', s_{-i}) = 1$.)

To complete the proof in general, suppose $m_1 > m_2 \geq m_3 \geq \cdots \geq m_n$. Let $u(s)$ be defined on $\mathcal{M} \cap \{s_1 \leq m_2\}$ as above. No matter how u is defined on the rest of $\mathcal{M}$, it will be sensitive to coordinate i for $2 \leq i \leq n$. In addition, the m_2 functions of s_{-1}, defined by $u(a, s_{-1})$ for $a \in \mathcal{M}_2$, are distinct binary valued functions of s_{-1}. Since m_1 is less than or equal to $2^{m_2 \cdots m_n}$, which is the total number of distinct binary valued functions of s_{-1}, it follows that u can be extended to the rest of $\mathcal{M}$ so that $u(a, s_{-1})$, for $1 \leq a \leq m_1$, are distinct binary valued functions of s_{-1}. The resulting function u is sensitive to each of its n coordinates. The lemma is proved. □.

Proof of Proposition 3.1 The mechanism used in the proof is a first price auction with a huge punishment for lying. Let $u : \mathcal{M} \to \{0,1\}$ be a function which is sensitive to every coordinate. Given the vector of reported values s', the seller compares b to $u(s')$. If they are not equal the seller imposes a huge fine on all the bidders. This means that if any bidder were to lie for any of its possible observed states, and if all other bidders were always truthful, then the seller would have a positive probability of detecting that someone has lied. This makes truth telling a Bayes-Nash equilibrium for the bidders. □

IV. On the Value of a Noisy Bit

Proposition 3.1 shows that a single bit can be a very effective piece of information. It relies rather heavily, however, on the bit being perfectly reliable. If there were a small chance that the value b were reported in error, then there would be a small chance that the seller would impose a huge fine even when the bidders were all truthful. Then truth telling would not be individually interim rational. This leads to the following question. The seller's side information b is said to be ϵ-reliable for some $\epsilon > 0$ if the function ϕ is restricted to $|\phi(s, b) - \frac{1}{2}| \leq \epsilon$ for all s and b. As ϵ approaches zero the side information becomes more nearly independent of the state vector s. Assuming the side information is ϵ reliable is equivalent to assuming that an arbitrary function $\tilde{\phi}$ is used to generate a bit $\tilde{b}$, and then b is generated by switching $\tilde{b}$ to $1 - \tilde{b}$ with probability $\frac{1}{2} - \epsilon$.

Proposition 4.1: Suppose $m_i \leq 1 + \prod_{j \neq i} m_j$ for each i. Then for any $\epsilon > 0$ there is a choice of ϵ- reliable side information for which there exists an FSE, IIR, BN-IC mechanism.

At first this result may seem surprising, but in light of the results of Crémer and McLean [1] it is easily understood. Crémer and McLean [1], building on an example of Myerson [6], show that if the states of the bidders are dependent, then, subject to a mild rank condition, FSE is possible. Arbitrarily small deviations from independence can still allow for the rank condition and therefore FSE to hold. Basically, even if the states of the bidders are independent, they can be conditionally dependent given b. The proof of the proposition, given next, is similar to the construction of [1].

Lemma 4.2: Suppose K and L are positive integers with $K \leq L + 1$. Consider the $K \times 2L$ random matrix U such that the variables $(U_{i,j} : 1 \leq i \leq K, 1 \leq j \leq L)$ are independent random variables, each continuously distributed over $[0, 1]$, and $U_{i,j+L} = 1 - U_{i,j}$ for $1 \leq i \leq K$ and $1 \leq j \leq L$). Then U has full rank (i.e. rank K) with probability one.

Proof Let e denote the K vector of all ones. Note that since $K - 1 \leq L$, the variables of the first $K - 1$ columns of U are independent, and uniformly distributed random variables. Thus, e and the first $K - 1$ columns of U are linearly independent with probability one. Equivalently, the first $K - 1$ columns and the $L + 1^{st}$ column of U are linearly independent with probability one. Thus, U has K linearly independent columns with probability one, and hence U has full rank with probability one. □

Proof of Proposition 4.1 The mechanism used in the proof will be a second price auction with a tax $h_i(s'_{-i}, b)$ imposed on bidder i for each i. In the absence of a tax, the interim expected gain of bidder i, assuming truth telling, is

$$g_i(s_i) = E_{s_{-i}}[(w_i(s_i) - \max_{j \neq i} w_j(s_j))_+ | s_i]$$

Since truth telling is a BN-IC strategy for a second price auction with no tax (in fact a dominant strategy for each bidder), and since the tax $h_i(s'_{-i}, b)$ to bidder i does not depend on the report of bidder i, the mechanism with tax is also BN-IC. (However, it could be that truth telling is not a dominant strategy for the second price auction with taxes.)

In order for the mechanism to be IIR and FSE it is necessary and sufficient that

$$g_i(s_i) = E[h_i(s_{-i}, b)|s_i] \tag{1}$$

For $1 \leq i \leq n$, $P[s_{-i}, b|s_i] = \pi_{-i}(s_{-i})\phi(s_i, s_{-i}, b)$.Thus

$$E[h_i(s_{-i}, b)|s_i] = \sum_{s_{-i}, b} h_i(s_{-i}, b)\pi_{-i}(s_{-i})\phi(s_i, s_{-i}, b).$$

For fixed i, the function h_i can be selected so that (1) is true for all values of s_i, as long as the $m_i \times (2\prod_{j \neq i} m_j)$ matrix with $s_i, (s_{-i}, b)^{th}$ entry equal to $\pi_{-i}(s_{-i})\phi(s_i, s_{-i}, b)$ has full row rank. This matrix is the product of two matrices: the first is the matrix of the same dimension with $s_i, (s_{-i}, b)^{th}$ entry equal to $\phi(s_i, s_{-i}, b)$, and the second matrix is a diagonal matrix with the nonzero terms of the form $\pi_{-i}(s_{-i})$ on the diagonal. Multiplication by the diagonal matrix does not change the row rank. Therefore, a sufficient condition for the existence of the taxes as required is that the $m_i \times (2\prod_{j \neq i} m_j)$ matrix with $s_i, (s_{-i}, b)^{th}$ entry equal to $\phi(s_i, s_{-i}, b)$ has full row rank. Suppose the values $\phi(s, b = 0)$ are independently distributed, each on the interval $[\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$ with some probability density. Then the matrices required to have full rank have the form as in Lemma 4.2. With probability one, the matrix has full rank. Since this is true for each i, it follows that the random choice of ϕ satisfies all the required rank conditions, with probability one. There thus exists an ϵ-reliable deterministic choice for ϕ satisfying the required rank conditions. □

V. Conclusion

In some sense, the results of this paper show that the way communications engineers and information theorists often measure information, either in bits or through the use of Shannon's notion of mutual information, may not be well suited for the modeling of limited communication among bidders. Or the result may simply be a reflection of a lack of robustness in the formulation of optimal auctions, indicated earlier by the example of Myerson [6] and results of Crémer and R. McLean [1]. Or the result my reflect the fact that even though only a single noisy bit is revealed to the seller, the seller has to use exact knowledge of the joint distribution of the single bit with the private valuations of the bidders.

The role of information, such as how much bidders know about the value of an object, and how much the seller knows about how much the bidders know about the value of an object, and so on, is central to much current research on mechanism design, based on Bayesian formulations. Perhaps the context of robust mechanisms for general environments (see, for example, [7]) would be a setting in which the notion of mutual information could be used to quantify the value of information in the context of auctions and mechanism design.

References

[1] J. Crémer and R. McLean, "Full extraction of the surplus in Bayesian and dominant strategy auctions," *Econometrica*, vol. 56, 1247-1258, 1988.

[2] D. Fudenberg and J. Tirole, *Game Theory*, MIT Press, Cambridge, MA, 2002.

[3] P. Klemperer, *Auctions; Theory and Practice*, Princeton University Press, 2004.

[4] V. Krishna, *Auction Theory*, Academic Press, 2002.

[5] E.S. Maskin and J.G. Riley, "Auction theory with private values," The American Economic Review, vol. 75, no. 2, pp. 150-155, 1985.

[6] R.B. Myerson, "Optimal auction design," *Mathematics of Operations Research,* vol. 6, no. 1, pp. 58-73, 1981.

[7] Z. Neeman, "The effectiveness of English auctions," *Games and Economic Behavior*, vol. 43, pp. 214-238, 2003.

[8] J.G. Riley and W.F. Samuelson, "Optimal auctions," *American Economic Review*, vol. 71, pp. 381-392, 1981.

[9] W. Vickrey, "Counter speculation, auctions, and competitive sealed tenders," *Journal of Finance*, vol. 16, no. 1, pp. 8-37, 1961.