

Cambridge University Press

0521818729 - Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume I

Edited by Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky

Excerpt

[More information](#)

CHAPTER 1

Auctions and Efficiency

Eric Maskin

1. INTRODUCTION

The allocation of resources is an all-pervasive theme in economics. Furthermore, the question of whether there exist mechanisms ensuring *efficient* allocation (i.e., mechanisms that ensure that resources end up in the hands of those who value them most) is of central importance in the discipline. Indeed, the very word “economics” connotes a preoccupation with the issue of efficiency.

But economists’ interest in efficiency does not end with the question of existence. If efficient mechanisms can be constructed, we want to know what they look like and to what extent they might resemble institutions used in practice.

Understandably, the question of what will constitute an efficient mechanism has been a major concern of economic theorists going back to Adam Smith. But, the issue is far from just a theoretical one. It is also of considerable practical importance. This is particularly clear when it comes to *privatization*, the transfer of assets from the state to the private sector.

In the last 15 years or so, we have seen a remarkable flurry of privatizations in Eastern Europe, the former Soviet Union, China, and highly industrialized Western nations, such as the United States, the United Kingdom, and Germany. An important justification for these transfers has been the expectation that they will improve efficiency. But if efficiency is the rationale, an obvious leading question to ask is: “What sorts of transfer mechanisms will *best* advance this objective?”

One possible and, of course, familiar answer is “the Market.” We know from the First Theorem of Welfare Economics (see Debreu, 1959) that, under certain conditions, the *competitive* mechanism (the uninhibited exchange and production of goods by buyers and sellers) results in an efficient allocation. A major constraint on the applicability of this result to the circumstances of privatization, however, is the theorem’s hypothesis of *large numbers*. For the competitive mechanism to work properly – to avoid the exercise of monopoly power – there must be sufficiently many buyers and sellers so that no single agent has an appreciable effect on prices. But privatization often entails small

Cambridge University Press

0521818729 - Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume I

Edited by Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky

Excerpt

[More information](#)

numbers. In the recent U.S. “spectrum” auctions – the auctions in which the government sold rights (in the form of licenses) to use certain radio frequency bands for telecommunications – there were often only two or three serious bidders for a given license. The competitive model does not seem readily applicable to such a setting.

An interesting alternative possibility was raised by William Vickrey (1961) 40 years ago. Vickrey showed that, if a seller has a single indivisible good for sale, a *second-price* auction (see Section 2) is an efficient mechanism – i.e., the winner is the buyer whose valuation of the good is highest – in the case where buyers have *private values* (“private values” mean that no buyer’s private information affects any other buyer’s valuation). This finding is rendered even more significant by the fact that it can be readily extended to the sale of multiple goods,¹ as shown by Theodore Groves (1973) and Edward Clarke (1971). Unfortunately, once the assumption of private values is dropped and thus buyers’ valuations *do* depend on other buyers’ information (i.e., we are in the world of common² or interdependent values), the second-price auction is no longer efficient, as I will illustrate later by means of an example. Yet, the common-values case is the norm in practice. If, say, a telecommunications firm undertakes a market survey to forecast demand for cell phones in a given region, the results of the survey will surely be of interest to its competitors and thus turn the situation into one of common values.

Recently, a literature has developed on the design of efficient auctions in common-values settings. The time is not yet ripe for a survey; the area is currently evolving too rapidly for that. But I would like to take this opportunity to discuss a few of the ideas from this literature.

2. THE BASIC MODEL

Because it is particularly simple, I will begin with the case of a single indivisible good. Later, I will argue that much (but not all) of what holds in the one-good case extends to multiple goods.

Suppose that there are n potential buyers. It will be simplest to assume that they are risk-neutral (however, we can accommodate any other attitude toward risk if the model is specialized to the case in which there is no residual uncertainty about valuations when all buyers’ information is pooled). Assume that each buyer i ’s private information about the good can be summarized by a *real-valued* signal. That is, buyer i ’s information is reducible to a one-dimensional parameter.³ Formally, suppose that each buyer i ’s signal s_i lies in

¹ Vickrey himself also treated the case of multiple units of the same good.

² I am using “common values” in the *broad* sense to cover any instance where one agent’s payoff depends on another’s information. The term is sometimes used narrowly to mean that all agents share the same payoff.

³ Later on, I will examine the case of multidimensional signals. As with multiple goods, much will generalize. As we will see, the most problematic case is that in which there are *both* multiple goods and multidimensional signals.

an interval $[\underline{s}_i, \bar{s}_i]$. The joint prior distribution of (s_1, \dots, s_n) is given by the c.d.f. $F(s_1, \dots, s_n)$. Buyer i 's valuation for the good (i.e., the most he would be willing to pay for it) is given by the function $v_i(s_1, \dots, s_n)$. I shall suppose (with little loss of generality) that higher values of s_i correspond to higher valuations, i.e.,

$$\frac{\partial v_i}{\partial s_i} > 0. \tag{2.1}$$

Let us examine two illustrations of this model.

Example 2.1. *Suppose that*

$$v_i(s_1, \dots, s_n) = s_i.$$

In this case, we are in the world of private values, not the interesting setting from the perspective of this lecture, but a valid special case.

A more pertinent example is:

Example 2.2. *Suppose that the true value of the good to buyer i is y_i , which, in turn, is the sum of a value component that is common to all buyers and a component that is peculiar to buyer i . That is,*

$$y_i = z + z_i,$$

where z is the common component and z_i is buyer i 's idiosyncratic component. Suppose, however, that buyer i does not actually observe y_i , but only a noisy signal

$$s_i = y_i + \varepsilon_i, \tag{2.2}$$

where ε_i is the noise term, and all the random variables $-z$, the z_i s, and the ε_i s—are independent. In this case, every buyer j 's signal s_j provides information to buyer i about his valuation, because s_j is correlated [via (2.2)] with the common component z . Hence, we can express $v_i(s_1, \dots, s_n)$ as

$$v_i(s_1, \dots, s_n) = E[y_i | s_1, \dots, s_n], \tag{2.3}$$

where the right-hand side of (2.3) denotes the expectation of y_i conditional on the signals (s_1, \dots, s_n) .

This second example might be kept in mind as representative of the sort of scenario that the analysis is intended to apply to.

3. AUCTIONS

An *auction* in the model of Section 2 is a *mechanism* (alternatively termed a “game form” or “outcome function”) that, on the basis of the bids submitted, determines (i) who wins (i.e., who – if anyone – is awarded the good), and

4 Maskin

(ii) how much each buyer pays.⁴ Let us call an auction *efficient* provided that, in equilibrium, buyer i is the winner if and only if

$$v_i(s_1, \dots, s_n) \geq \max_{j \neq i} v_j(s_1, \dots, s_n) \quad (3.1)$$

(this definition is slightly inaccurate because of the possibility of ties for highest valuation, an issue that I shall ignore). In other words, efficiency demands that, in an equilibrium of the auction, the winner be the buyer with the highest valuation, conditional on *all available information* (i.e., on all buyers' signals).

This notion of efficiency is sometimes called *expost efficiency*. It assumes implicitly that the social value of the good being sold equals the maximum of the potential buyers' individual valuations. This assumption would be justified if, for example, each buyer used the good (e.g., a spectrum license) to produce an output (e.g., telecommunication service) that is sold in a competitive market without significant externalities (market power or externalities might drive a wedge between individual and social values).

The reader may wonder why, even if one wants efficiency, it is necessary to insist that the auction itself be efficient. After all, the buyers could always retrade afterward if the auction resulted in a winner with less than the highest valuation. The problem with relying on postauction trade, however, is much the same as that plaguing competitive exchange in the first place: These mechanisms do not, in general, work efficiently when there are only a few traders. To see this, consider the following example:⁵

Example 3.1. *Suppose that there are two buyers. Assume that buyer 1 has won the auction and has a valuation of 1. If the auction is not guaranteed to be efficient, then there is some chance that buyer 2's valuation is higher. Suppose that, from buyer 1's perspective, buyer 2's valuation is distributed uniformly in the interval $[0, 2]$. Now, if there is to be further trade after the auction, someone has to initiate it. Let us assume that buyer 1 does so by proposing a trading price to buyer 2. Presumably, buyer 1 will propose a price p^* that maximizes his expected payoff, i.e., that solves*

$$\max_p \frac{1}{2}(2 - p)(p - 1). \quad (*)$$

[To understand (*), note that $\frac{1}{2}(2 - p)$ is the probability that the proposal is accepted – since it is the probability that buyer 2's valuation is at least p – and that $p - 1$ is buyer 1's net gain in the event of acceptance.] But the solution to (*) is $p^* = \frac{3}{2}$. Hence, if buyer 2's valuation lies between 1 and $\frac{3}{2}$, the allocation,

⁴ For some purposes – e.g., dealing with risk-averse buyers (see Maskin and Riley, 1984), liquidity constraints (see Che and Gale, 1996, or Maskin, 2000) or allocative externalities (see Jehiel and Moldovanu (2001) – one must consider auctions in which buyers other than the winner also make payments. In this lecture, however, I will not have to deal with this possibility.

⁵ In this example, buyers have private values, but, as Fieseler, Kittsteiner, and Moldovanu (2000) show, resale can become even more problematic when there are common values.

Cambridge University Press

0521818729 - Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume I

Edited by Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky

Excerpt

[More information](#)

Auctions and Efficiency

5

even after allowing for ex post trade, will remain inefficient, because buyer 2 will reject 1's proposal.

I will first look at efficiency in the *second-price* auction. This auction form (often called the *Vickrey* auction) has the following rules: (i) each bidder i makes a (sealed) bid b_i , which is a nonnegative number; (ii) the winner is the bidder who has made the highest bid (again ignoring the issue of ties); (iii) the winner pays the second-highest bid, $\max_{j \neq i} b_j$. As I have already noted and will illustrate explicitly, in Section 6 this auction can readily be extended to multiple goods.

The Vickrey auction is efficient in the case of private values.⁶ To see this, note first that it is optimal – in fact, a dominant strategy – for buyer i to set $b_i = v_i$ (i.e., to bid his true valuation). In particular, bidding below v_i does not affect buyer i 's payment if he wins (because his bid does not depend on his own bid); it just reduces his chance of winning – and so is not a good strategy. Bidding above v_i raises buyer i 's probability of winning, but the additional events in which he wins are precisely those in which someone else has bid higher than v_i . In such events, buyer i pays more than v_i , also not a desirable outcome. Thus, it is indeed optimal to bid $b_i = v_i$, which implies that the winner is the buyer with the highest valuation, the criterion for efficiency.

Unfortunately, the Vickrey auction does not remain efficient once we depart from private values. To see this, consider the following example.

Example 3.4. Suppose that there are three buyers with valuation functions

$$v_1(s_1, s_2, s_3) = s_1 + \frac{2}{3}s_2 + \frac{1}{3}s_3,$$

$$v_2(s_1, s_2, s_3) = s_2 + \frac{1}{3}s_1 + \frac{2}{3}s_3,$$

$$v_3(s_1, s_2, s_3) = s_3.$$

Notice that buyers 1 and 2 have common values (i.e., their valuations do not depend only on their own signals). Assume that it happens that $s_1 = s_2 = 1$ (of course, buyers 1 and 2 would not know that their signal values are equal, because signals are private information), and suppose that buyer 3's signal value is either slightly below or slightly above 1. In the former case, it is easy to see that

$$v_1 > v_2 > v_3,$$

and so, for efficiency, buyer 1 ought to win. However, in the latter case

$$v_2 > v_1 > v_3,$$

⁶ It is easy to show that the “first-price” auction – the auction in which each buyer makes a bid, the high bidder wins, and the winner pays his bid – is a nonstarter as far as efficiency is concerned. Indeed, even in the case of private values, the first-price auction is never efficient, except when buyers' valuations are symmetrically distributed (see Maskin, 1992).

and so buyer 2 is the efficient winner. Thus, the efficient allocation between buyers 1 and 2 turns on whether s_3 is below or above 1. But, in a Vickrey auction, the bids made by buyers 1 and 2 cannot incorporate information about s_3 , because that signal is private information to buyer 3. Thus, the outcome of the auction cannot in general be efficient.

4. AN EFFICIENT AUCTION

How should we respond to the shortcomings of the Vickrey auction as illustrated by Example 3.3? One possible reaction is to appeal to classical mechanism-design theory. Specifically, we could have each buyer i announce a signal value \hat{s}_i , award the good to the buyer i for whom $v_i(\hat{s}_1, \dots, \hat{s}_n)$ is highest, and choose the winner's payment to evoke truth-telling in buyers (i.e., to induce each buyer j to set \hat{s}_j equal to his true signal value s_j). This approach is taken in Crémer and McLean (1985) and Maskin (1992).

The problem with such a “direct revelation” mechanism is that it is utterly unworkable in practice. In particular, notice that it requires the mechanism designer to know the physical signal spaces S_1, \dots, S_n , the functional forms $v_i(\cdot)$, and the prior distributions of the signals – an extraordinarily demanding constraint. Now, the mechanism designer could attempt to elicit this information from the buyers themselves using the methods of the implementation literature (see Palfrey, 1993). For example, to learn the signal spaces, he could have each buyer announce a vector $(\hat{S}_1, \dots, \hat{S}_n)$ and assign suitable penalties if the announcements did not match up appropriately. A major difficulty with such a scheme, however, is that in all likelihood the signal spaces S_i are themselves *private* information. For analytic purposes, we model S_i as simply an interval of numbers. But, this abstracts from the reality that buyer i 's signal corresponds to some *physical* entity – whatever it is that buyer i observes. Indeed, the signal may well be a sufficient statistic for data from a variety of different informational sources, and there is no reason why other buyers should know just what this array of sources is.

To avoid these complications, I shall concentrate on auction rules that do not make use of such details as signal spaces, functional forms, and distributions. Indeed, I will be interested in auctions that work well irrespective of these details; that is, I will adhere to the “Wilson Doctrine” (after Robert Wilson, who has been an eloquent proponent of the view that auction institutions should be “detail-free”). It turns out that a judicious modification of the Vickrey auction will do the trick.

Before turning to the modification, however, I need to introduce a restriction on valuation functions that is critical to the possibility of constructing efficient auctions. Let us assume that for all i and $j \neq i$ and all (s_1, \dots, s_n) ,

$$v_i(s_1, \dots, s_n) = v_j(s_1, \dots, s_n) \Rightarrow \frac{\partial v_i}{\partial s_i}(s_1, \dots, s_n) > \frac{\partial v_j}{\partial s_i}(s_1, \dots, s_n).^7 \quad (4.1)$$

⁷ This condition was introduced by Gresik (1991).

In other words, condition (4.1) says that buyer i 's signal has a greater marginal effect on his own valuation than on that of any other buyer j (at least at points where buyer i 's and buyer j 's valuations are equal).

Notice that, in view of (2.1), condition (4.1)⁸ is automatically satisfied by Example 2.1 (the case of private values): the right-hand side of the inequality then simply vanishes. Condition (4.1) also holds for Example 2.2. This is because, in that example, s_i conveys relevant information to buyer j ($\neq i$) about the common component z , but tells buyer i not only about z but also his idiosyncratic component z_i . Thus, v_i will be more sensitive than v_j to variations in s_i .

But whether or not condition (4.1) is likely to be satisfied, it is, in any event, essential for efficiency. To see what can go wrong without it, consider the following example.

Example 4.5. Suppose that the owner of a tract of land wishes to sell off the rights to drill for oil on her property. There are two potential drillers who are competing for this right. Driller 1's fixed cost of drilling is 1, whereas his marginal cost is 2. In contrast, driller 2 has fixed and marginal costs of 2 and 1, respectively. Assume that driller 1 observes how much oil is underground. That is, s_1 equals the quantity of oil. Driller 2 obtains no private information. Then, if the price of oil is 4, we have

$$\begin{aligned} v_1(s_1) &= (4 - 2)s_1 - 1 = 2s_1 - 1, \\ v_2(s_1) &= (4 - 1)s_1 - 2 = 3s_1 - 2. \end{aligned}$$

Observe that $v_1(s_1) > v_2(s_1)$ if and only if $s_1 < 1$. Thus, for efficiency, driller 1 should be awarded drilling rights provided that $\frac{1}{2} < s_1 < 1$ (for $s_1 < \frac{1}{2}$, there is not enough oil to justify drilling at all). Driller 2, by contrast, should get the rights when $s_1 > 1$.

In this example, there is no way (either through a modified Vickrey auction or otherwise) of inducing driller 1 to reveal the true value s_1 to allocate drilling rights efficiently. To see this, consider, without loss of generality, a direct revelation mechanism and let $t_1(\hat{s}_1)$ be a monetary transfer (possibly negative) to driller 1 if he announces signal value \hat{s}_1 . Let s'_1 and s''_1 be signal values such that

$$\frac{1}{2} < s'_1 < 1 < s''_1. \tag{4.2}$$

Then, for driller 1 to have the incentive to announce truthfully when $s_1 = s''_1$, we must have

$$t_1(s''_1) \geq 2s''_1 - 1 + t_1(s'_1) \tag{4.3}$$

⁸ Notice that the strictness of the inequality in (4.1) rules out the case of "pure common values," where all buyers share the same valuation. However, in that case, the issue of who wins does not matter for efficiency.

(the left-hand side is his payoff when he is truthful, whereas the right-hand side is his payoff when he pretends that $s_1 = s'_1$). Similarly, the incentive-constraint corresponding to $s_1 = s'_1$ is

$$2s'_1 - 1 + t_1(s'_1) \geq t_1(s''_1). \quad (4.4)$$

Subtracting (4.4) from (4.3), we obtain

$$2(s'_1 - s''_1) \geq 0,$$

a contradiction of (4.2). Hence, there exists no efficient mechanism.

The feature that interferes with efficiency in this example is the violation of condition (4.1), i.e., the fact that

$$0 < \frac{\partial v_1}{\partial s_1} < \frac{\partial v_2}{\partial s_1}. \quad (4.5)$$

Inequalities (2.1) and (4.5) imply that, as s_1 rises, drilling rights become more valuable to driller 1, but increasingly more likely, from the standpoint of efficiency, to be awarded to driller 2. This conflict makes the task of providing proper incentives for driller 1 impossible.

Assuming henceforth that (4.1) holds, let us reconfront the task of designing an efficient auction. In Example 3.4, we saw that the Vickrey auction failed because buyers 1 and 2 could not incorporate pertinent information about buyer 3 in their bids (since s_3 was private information). This suggests that, as in Dasgupta and Maskin (2000), a natural way of amending the Vickrey auction would be to allow buyers to make *contingent* bids – bids that depend on other buyers' valuations. In Example 3.4, this would enable buyer 1 to say, in effect, "I don't know what buyer 3's valuation is, but if it turns out to be x , then I want to bid y ."

Let us examine how contingent bidding would work in the case of two buyers. Buyer 1 would announce a schedule $\hat{b}_1(\cdot)$, where, for all possible values v_2 ,

$$\hat{b}_1(v_2) = \text{buyer 1's bid if buyer 2 has valuation } v_2.$$

Similarly, buyer 2 would announce a schedule $\hat{b}_2(\cdot)$, where

$$\hat{b}_2(v_1) = \text{buyer 2's bid if buyer 1's valuation is } v_1.$$

We would then look for a fixed point

$$(v_1^o, v_2^o) = (\hat{b}_1(v_2^o), \hat{b}_2(v_1^o)) \quad (4.6)$$

and

$$\text{install buyer 1 as the winner if and only if } v_1^o > v_2^o. \quad (4.7)$$

To understand the rationale for (4.6) and (4.7), imagine that buyers bid *truthfully*. Because signals are private information and thus buyer 1 will not in general know his own valuation, truthful bidding means that, if his signal value

is s_1 , he submits a schedule $\hat{b}_1(\cdot) = b_1(\cdot)$ such that

$$b_1(v_2(s_1, s'_2)) = v_1(s_1, s'_2) \quad \text{for all } s'_2. \quad (4.8)$$

That is, whatever s'_2 (and hence v_2) turns out to be, buyer 1 bids his true valuation for that signal value. Similarly, truthful bidding for buyer 2 with signal value s_2 means reporting schedule $\hat{b}_2(\cdot) = b_2(\cdot)$, such that

$$b_2(v_1(s'_1, s_2)) = v_2(s'_1, s_2) \quad \text{for all } s'_1. \quad (4.9)$$

Observe that if buyers bid according to (4.8) and (4.9), then the true valuations

$$(v_1(s_1, s_2), v_2(s_1, s_2))$$

constitute a fixed point in the sense of (4.6).¹⁰

In view of (4.6) and (4.7), this means that if buyers are truthful, the auction will result in an efficient allocation. Thus, the remaining critical issue is how to get buyers to bid truthfully. For this purpose, it is useful to recall the device that the Vickrey auction exploits to induce truthful bidding, viz. to make the winner's payment equal, not to his own bid, but to the lowest possible bid he could have made and still have won the auction.

This trick cannot be exactly replicated in our setting because buyers are submitting schedules rather than single bids. But let us try to take it as far as it will go. Suppose that when buyers report the schedules $(\hat{b}_1(\cdot), \hat{b}_2(\cdot))$, the resulting fixed point (v_1^o, v_2^o) satisfies

$$v_1^o > v_2^o.$$

Then, according to our rules, buyer 1 should win. But rather than having him pay v_1^o , we will have buyer 1 pay v_1^* , where

$$v_1^* = \hat{b}_2(v_1^*). \quad (4.10)$$

This payment rule, I maintain, is the common-values analog of the Vickrey trick in the sense that v_1^* is the lowest *constant* bid (i.e., the lowest uncontingent bid) that buyer 1 could make and still win (or tie for winning) given buyer 2's bid $\hat{b}_2(\cdot)$. The corresponding payment rule for buyer 2 should he win is v_2^* such that

$$v_2^* = \hat{b}_1(v_2^*). \quad (4.11)$$

I claim that, given the payment rules (4.10) and (4.11), it is an equilibrium for buyers to bid truthfully. To see this most easily, let us make use of a strengthened

⁹ I noted in my arguments against direct revelation mechanisms that buyer 1 most likely will not know buyer 2's signal space S_2 . But this in no way should prevent him from understanding how his own valuation is related to that of buyer 2, which is what (4.8) is *really* expressing [i.e., (4.8) still makes sense even if buyer 1 does not know what values s'_2 can take].

¹⁰ Without further assumptions on valuation functions, there could be additional – nontruthful – fixed points. Dasgupta and Maskin (2000) and Eso and Maskin (2000a) provide conditions to rule such fixed points out. But even if they are not ruled out, the auction rules can be modified so that, in equilibrium, the truthful fixed point results (see Dasgupta and Maskin, 2000).

10 **Maskin**

version of (4.1):

$$\frac{\partial v_i}{\partial s_i} > \frac{\partial v_j}{\partial s_i}. \quad (4.12)$$

Let us suppose that buyer 2 is truthful, i.e., he bids $b_2(\cdot)$ satisfying (4.9). I must show that it is optimal for buyer 1 to bid $b_1(\cdot)$ satisfying (4.8).

Notice first that if buyer 1 wins, his payoff is

$$v_1(s_1, s_2) - v_1^*, \quad \text{where} \quad v_1^* = b_2(v_1^*), \quad (4.13)$$

regardless of how he bids (because neither his valuation nor his payment depends on his bid). I claim that if buyer 1 bids truthfully, then he wins if and only if (4.13) is positive. Observe that if this claim is established, then I will in fact have shown that truthful bidding is optimal; because buyer 1's bid does not affect (4.13), the most he can possibly hope for is to win precisely in those cases where the net payoff from winning is positive.

To see that the claim holds, let us first differentiate (4.9) with respect to s'_1 to obtain

$$\frac{db_2}{dv_1}(v_1(s'_1, s_2)) \frac{\partial v_1}{\partial s_1}(s'_1, s_2) = \frac{\partial v_2}{\partial s_1}(s'_1, s_2) \quad \text{for all} \quad s'_1.$$

This identity, together with (2.1) and (4.12), implies that

$$\frac{db_2}{dv_1}(v_1) < 1, \quad \text{for all} \quad v_1. \quad (4.14)$$

But, from (4.14), (4.13) is positive if and only if

$$v_1(s_1, s_2) - v_1^* > \frac{db_2}{dv_1}(v'_1)(v_1(s_1, s_2) - v_1^*) \quad \text{for all} \quad v'_1. \quad (4.15)$$

Now, from the intermediate value theorem, there exists $v'_1 \in [v_1^*, v_1(s_1, s_2)]$ such that

$$b_2(v_1(s_1, s_2)) - b_2(v_1^*) = \frac{db_2}{dv_1}(v'_1)(v_1(s_1, s_2) - v_1^*).$$

Hence (4.13) is positive if and only if

$$v_1(s_1, s_2) - v_1^* > b_2(v_1(s_1, s_2)) - b_2(v_1^*), \quad (4.16)$$

which, because $v_1^* = b_2(v_1^*)$, is equivalent to

$$v_1(s_1, s_2) > v_2(s_1, s_2). \quad (4.17)$$

Now suppose that buyer 1 is truthful. Because $(v_1(s_1, s_2), v_2(s_1, s_2))$ is then a fixed point, 1 wins if and only if (4.17) holds. So, we can conclude that, when buyer 1 is truthful, his net payoff from winning is positive [i.e., (4.13) is positive] if and only if he wins, which is what I claimed. That is, the modified Vickrey auction is efficient. (This analysis ignores the possible costs to buyers of acquiring signals; once such costs are incorporated the modified Vickrey