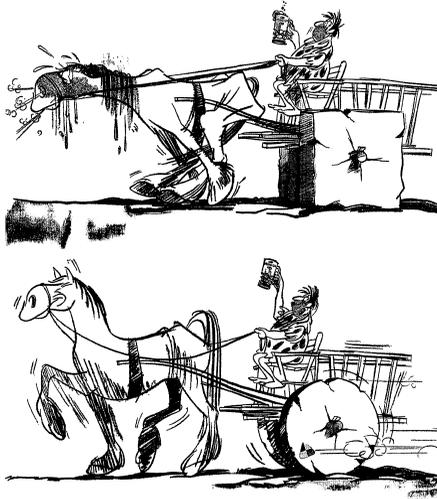


Why Every Economist Should Learn Some Auction Theory*

We discuss the strong connections between auction theory and “standard” economic theory; we show that situations that do not at first sight look like auctions can be recast to use auction-theoretic techniques; and we argue that auction-theoretic tools and intuitions can provide useful arguments and insights in a broad range of mainstream economic settings. We also discuss some more obvious applications, especially to industrial organization.



Auction Theory **Heineken.**
Refreshes the parts other
economics **beers cannot reach**

Disclaimer: We do not contend that the following ideas are all as important as the one illustrated here, merely that those who have not imbibed auction theory are missing out on a potent brew!

* This was an Invited Lecture to the 8th World Congress of the Econometric Society, and was originally published under the title “Why Every Economist Should Learn Some Auction Theory”, in Dewatripont, M., Hansen, L., and Turnovsky, S. (eds.) (2003) *Advances in Economics and Econometrics: Theory and Applications*. Cambridge University Press, vol. 1, pp. 25–55. Susan Athey was an excellent discussant of my lecture. I also received extremely helpful comments and advice from many other friends and colleagues, including Larry Ausubel, Mike Baye, Alan Beggs, Simon Board, Jeremy Bulow, Peter Cramton, Joe Farrell, Giulio Federico, Nils Hendrik von der Fehr, Dan Kovenock, David McAdams, Peter McAfee, Flavio Menezes, Meg Meyer, Jonathan

2.1 INTRODUCTION

Auction theory has attracted enormous attention in the last few years.¹ It has been increasingly applied in practice, and this has itself generated a new burst of theory. It has also been extensively used, both experimentally and empirically, as a testing ground for game theory.² Furthermore, by carefully analyzing very simple trading models, auction theory is developing the fundamental building blocks for our understanding of more complex environments. But some people still see auction theory as a rather specialized field, distinct from the main body of economic theory, and as an endeavor for management scientists and operations researchers rather than as a part of mainstream economics. This chapter aims to counter that view.

This view may have arisen in part because auction theory was substantially developed by operational researchers, or in operations research journals,³ and using technical mathematical arguments rather than standard economic intuitions. But it need not have been this way. This chapter argues that the connections between auction theory and “standard” economic theory run deeper than many people realize; that auction-theoretic tools provide useful arguments in a broad range of contexts; and that a good understanding of auction theory is valuable in developing intuitions and insights that can inform the analysis of many mainstream economic questions. In short, auction theory is central to economics.

We pursue this agenda in the context of some of the main themes of auction theory: the revenue equivalence theorem, marginal revenues, and ascending vs. (first-price) sealed-bid auctions. To show how auction-theoretic tools can be applied elsewhere in economics, section 2.2 exploits the revenue equivalence theorem to analyze a wide range of applications that are not, at first sight, auctions, including litigation systems, financial crashes, queues, and wars of attrition. (Section 4.6 briefly discusses an application of the theory of affiliation.) To illustrate how looser analogies can usefully be made between auction theory and economics, section 2.3 applies some intuitions from the compar-

Mirrlees-Black, John Morgan, Marco Pagnozzi, Nicola Persico, Eric Rasmussen, David Salant, Margaret Stevens, Rebecca Stone, Lucy White, Mark Williams, Xavier Vives, Caspar de Vries, and Charles Zheng.

¹ See chapter 1 for a review of auction theory; many of the most important contributions are collected in Klemperer (2000a).

² Kagel (1995) and Laffont (1997) are excellent recent surveys of the experimental and empirical work, respectively. Section 2.6 and chapters 3–7 discuss practical applications.

³ The earliest studies appear in the operations research literature, for example, Friedman (1956). Myerson’s (1981) breakthrough article appeared in *Mathematics of Operations Research*, while Rothkopf’s (1969) and Wilson’s (1967, 1969) classic early papers appeared in *Management Science*. Ortega Reichert’s (1968a) pathbreaking models of auctions, including a model of signaling that significantly predated Spence (1972), remain relatively little known by economists, perhaps because they formed an operations research PhD thesis.

ison of ascending and sealed-bid auctions to other economic settings such as rationing, and e-commerce. To demonstrate the deeper connections between auction theory and economics, section 2.4 discusses and applies the close parallel between the optimal auction problem and that of the discriminating monopolist; both are about maximizing marginal revenues.

Furthermore, auction-theoretic ways of thinking are also underutilized in more obvious areas of application, for instance, price-setting oligopolies which we discuss in section 2.5.⁴ Few non-auction-theorists know, for example, that marginal-cost pricing is *not* always the only equilibrium when identical firms with constant marginal costs set prices, or know the interesting implications of this fact. Section 2.6 briefly discusses direct applications of auction theory to markets that are literally auction markets, including electricity markets, treasury auctions, spectrum auctions, and internet markets, and we conclude in section 2.7.

2.2 USING AUCTION-THEORETIC TOOLS IN ECONOMICS: THE REVENUE EQUIVALENCE THEOREM

Auction theory's most celebrated theorem, the revenue equivalence theorem (RET) states conditions under which different auction forms yield the same expected revenue, and also allows revenue rankings of auctions to be developed when these conditions are violated.⁵ Our purpose here, however, is to apply it in contexts where the use of an auction model might not seem obvious.

Revenue Equivalence Theorem (RET). *Assume each of a given number of risk-neutral potential buyers has a privately known valuation independently drawn from a strictly increasing atomless distribution, and that no buyer wants more than one of the k identical indivisible prizes. Then any mechanism in which (i) the prizes always go to the k buyers with the highest valuations and (ii) any bidder with the lowest feasible valuation expects zero surplus, yields the same expected revenue (and results in each bidder making the same expected payment as a function of her valuation).*⁶

⁴ Of course, standard auction models form the basic building blocks of models in many contexts. See, for example, Stevens' (1994, 2000) models of wage determination in oligopsonistic labor markets, Bernheim and Whinston (1986), Feddersen and Pesendorfer (1996, 1998), Persico (2000a), and many others' political economy models, and many models in finance (including, of course, takeover battles, to which we give an application in section 2.4).

Another major area we do not develop here is the application of auction-theorists' understanding of the winner's curse to adverse selection more generally.

⁵ See chapter 1.

⁶ See Appendix 1.A for more general statements and an elementary proof. The theorem was first derived in an elementary form by Vickrey (1961, 1962) and subsequently extended to greater generality by Myerson (1981), Riley and Samuelson (1981), and others.

More general statements are possible, but are not needed for the current purpose. Our first example is very close to a pure auction.

2.2.1 Comparing Litigation Systems

In 1991, US Vice President Dan Quayle suggested reforming the US legal system in the hope, in particular, of reducing legal expenditures. One of his proposals was to augment the current rule according to which parties pay their own legal expenses, by a rule requiring the losing party to pay the winner an amount equal to the loser's own expenses. Quayle's intuition was that if spending an extra \$1 on a lawsuit might end up costing you \$2, then less would be spent. Was he correct?⁷

A simple starting point is to assume each party has a privately known value of winning the lawsuit relative to losing, independently drawn from a common, strictly increasing, atomless distribution;⁸ that the parties independently and simultaneously choose how much money to spend on legal expenses; and that the party who spends the most money wins the "prize" (the lawsuit).⁹ It is not too hard to see that both the existing US system and the Quayle system satisfy the assumptions of the RET, so the two systems result in the same expected total payments on lawyers.¹⁰ So Quayle was wrong (as usual); his argument is precisely offset by the fact that the value of winning the lawsuit is greater when you win your opponent's expenses.¹¹

Ah, Quayle might say, but this calculation has taken as given the set of lawsuits that are contested. Introducing the Quayle scheme will change the

⁷ This question was raised and analyzed (though not by invoking the RET) by Baye, Kovenock, and de Vries (1997). The ideas in this section, except for the method of analysis, are drawn from them. See also Baye, Kovenock, and de Vries (1998). (See chapter 1, exercise 4.)

⁸ For example, a suit about which party has the right to a patent might fit this model. The results extend easily to common-value settings, for example, contexts in which the issue is the amount of damages that should be transferred from one party to another.

⁹ American seminar audiences typically think this is a natural assumption, but non-Americans often regard it as unduly jaundiced. Of course, we use it as a benchmark only, to develop insight and intuition (just as the lowest price does not win the whole market in most real "Bertrand" markets, but making the extreme assumption is a common and useful starting point). Extensions are possible to cases in which with probability $(1 - \lambda)$ the "most deserving" party wins, but with probability $\lambda > 0$ the biggest spender wins.

¹⁰ The fact that no single "auctioneer" collects the players' payments as revenues, but that they are instead dissipated in legal expenses in competing for the single available prize (victory in the lawsuit), is of course irrelevant to the result.

Formally, checking our claims requires confirming that there are equilibria of the games that satisfy the RET's assumptions. The assumption we made that the parties make a one-shot choice of legal expenses is not necessary but makes confirming this relatively easy. See Baye, Kovenock, and de Vries (1997) for explicit solutions.

¹¹ Some readers might argue they could have inferred the effectiveness of the proposal from the name of the proponent, without need of further analysis. In fact, however, this was one of Dan Quayle's policy interventions that was not subject to immediate popular derision.

“bidding functions”, that is, change the amount any given party spends on litigation, so also change who decides to bring suits. Wrong again Dan! Although it is correct that the bidding functions change, the RET also tells us (in its parenthetical remark) that any given party’s *expected* payoffs from the lawsuit are unchanged, so the incentives to bring lawsuits are unchanged.

What about other systems, such as the typical European system in which the loser pays a fraction of the winner’s expenses? This is a trick question: it is no longer true that a party with the lowest possible valuation can spend nothing and lose nothing. Now this party always loses in equilibrium and must pay a fraction of the winner’s expenses, so makes negative expected surplus. That is, condition (ii) of the RET now fails. Thinking through the logic of the proof of the RET makes clear that all the players are worse off than under the previous systems.¹² That is, legal bills are higher under the European rule. The reason is that the incentives to win are greater than in the US system, and there is no offsetting effect. Here of course the issue of who brings lawsuits is important since low-valuation parties would do better not to contest suits in this kind of system; consistent with our theory there is empirical evidence (e.g., Hughes and Snyder, 1995) that the American system leads to more trials than, for example, the British system.

This last extension demonstrates that even where the RET in its simplest form fails, it is often possible to see how the result is modified; Appendix 2.A shows how to use the RET to solve for the relative merits of a much broader class of systems of which those we have discussed are special cases. We also show there that a system that might be thought of as the exact opposite of Quayle’s system is optimal in this model. Of course, many factors are ignored (e.g., asymmetries); the basic model should be regarded as no more than a starting point for analysis.

2.2.2 *The War of Attrition*

Consider a war of attrition in which N players compete for a prize. For example, N firms compete to be the unique survivor in a natural monopoly market, or N firms each hold out for the industry to adopt the standard they prefer.¹³ Each player pays costs of 1 per unit time until she quits the game. When just one player remains, that player also stops paying costs and wins the prize. There is no discounting. The two-player case, where just one quit is

¹² As Appendix 2.A discusses, every type’s surplus is determined by reference to the lowest-valuation type’s surplus (see also Appendix 1.A), and the lowest type is worse off in the European system. Again, our argument depends on condition (i) of the RET applying. See Appendix 2.A and Baye, Kovenock, and de Vries (1997).

¹³ Another related example analyzed by Bulow and Klemperer (1999) is that of N politicians each delaying in the hope of being able to avoid publicly supporting a necessary but unpopular policy that requires the support of $N - 1$ to be adopted.

needed to end the game, has been well analyzed.¹⁴ Does the many-player case yield anything of additional interest?

Assume players' values of winning are independently drawn from a common, strictly increasing, atomless distribution, and the game has an equilibrium satisfying the other conditions of the RET. Then the RET tells us that in expectation the total resources spent by the players in the war of attrition equal those paid by the players in any other mechanism satisfying the RET's conditions—for example, a standard ascending auction in which the price rises continuously until just one player remains and (only) the winner pays the final price. This final price will equal the second-highest actual valuation, so the expected total resources dissipated in the war of attrition is the expectation of this quantity.

Now imagine the war of attrition has been under way long enough that just the two highest-valuation players remain. What are the expected resources that will be dissipated by the remaining two players, starting from this time on? The RET tells us that they equal the auctioneer's expected revenue if the war of attrition were halted at this point and the objects sold to the remaining players by an ascending auction, that is, the expected second-highest valuation of these two remaining players. This is the same quantity, on average, as before!¹⁵ So the expected resources dissipated, and hence the total time taken until just two players remain, must be zero; all but the two highest-valuation players must have quit at once.

Of course this conclusion is, strictly speaking, impossible; the lowest-valuation players cannot identify who they are in zero time. However, the conclusion is correct in spirit, in that it is the limit point of the unique symmetric equilibria of a sequence of games which approach this game arbitrarily closely (and there is no symmetric equilibrium of the limit game).¹⁶

¹⁴ See, for example, Maynard Smith (1974) and Riley (1980) who discuss biological competition, Fudenberg and Tirole (1986) who discuss industrial competition, Abreu and Gul (2000), Kambe (1999), and others who analyze bargaining, and Bliss and Nalebuff (1984) who give a variety of amusing examples.

Bliss and Nalebuff note that extending to $K + 1$ players competing for K prizes does not change the analysis in any important way, since it remains true that just one quit is needed to end the game.

¹⁵ Of course the expectation of the second-highest valuation of the last two players is computed when just these two players remain, rather than at the beginning of the war of attrition as before. But on average these two expectations must be the same.

¹⁶ Bulow and Klemperer (1999) analyze games in which each player pays costs at rate 1 before quitting but must continue to pay costs even after quitting at rate c per unit time until the whole game ends. The limit $c \rightarrow 0$ corresponds to the war of attrition discussed here. (The case $c = 1$ corresponds, for example, to "standards battles" or political negotiations in which all players bear costs equally until *all* have agreed on the same standard or outcome; this game also has interesting properties—see Bulow and Klemperer.) Other series of games, for example, games in which being k th to last to quit earns a prize of ε^{k-1} times one's valuation, with $\varepsilon \rightarrow 0$, or games in which players can only quit at the discrete times $0, \varepsilon, 2\varepsilon, \dots$, with $\varepsilon \rightarrow 0$, also yield the same outcome in the limit.

Here, therefore, the role of the RET is less to perform the ultimate analysis than it is to show that there is an interesting and simple result to be obtained.¹⁷ Of course by developing intuition about what the result must be, the RET also makes proving it much easier. Furthermore the RET was also useful in the actual analysis of the more complex games that Bulow and Klemperer (1999) used to approximate this game. In addition, anyone armed with a knowledge of the RET can simplify the analysis of the basic two-player war of attrition.

2.2.3 *Queuing and Other “All-Pay” Applications*

The preceding applications have both been variants of “all-pay” auctions. As another elementary example of this kind consider different queuing systems, for example, for tickets to a sporting event. Under not unreasonable assumptions, a variety of different rules of queue management, for example, making the queue more or less comfortable, informing or not informing people whether the number queuing exceeds the number who will receive a ticket, etc., will make no difference to the social cost of the queuing mechanism. As in our litigation example (section 2.2.1), we think of these results as a starting point for analysis rather than as final conclusions.¹⁸

Many other issues such as lobbying battles, political campaigns,¹⁹ tournaments in firms, contributions to public goods,²⁰ patent races, and some kinds of price-setting oligopoly (see section 2.5.2) can be modeled as all-pay auctions and may provide similar applications.

2.2.4 *Solving for Equilibrium Behavior: Market Crashes and Trading “Frenzies”*

The examples thus far have all proceeded by computing the expected total payments made by all players. But the RET also states that each individual’s

¹⁷ It was the RET that showed Bulow and Klemperer that there was an analysis worth doing. Many people, and some literature, had assumed the many-player case would look like the two-player case but with more-complicated expressions, although Fudenberg and Kreps (1987) and Haigh and Cannings (1989) observed a similar result to ours in games without any private information and in which all players’ values are equal.

However, an alternative way to see the result in our war of attrition is to imagine the converse but that a player is within ε of her planned quit time when $n > 1$ other players remain. Then the player’s cost of waiting as planned is of order ε , but her benefit is of order ε^n since only when all n other players are within ε of giving up will she ultimately win. So for small ε she will prefer to quit now rather than wait, but in this case she should of course have quit ε earlier, and so on. So only when $n = 1$ is delay possible.

¹⁸ Holt and Sherman (1982) compute equilibrium behavior and hence obtain these results without using the RET.

¹⁹ See, especially, Persico (2000a).

²⁰ Menezes, Monteiro, and Temimi (2000) uses the RET in this context.

expected payment must be equal across mechanisms satisfying the assumptions. This fact can be used to infer what players' equilibrium actions must be in games which would be too complex to solve by any direct method of computing optimal behavior.²¹

Consider the following model. The aim is to represent, for example, a financial or housing market and show that trading "frenzies" and price "crashes" are the inevitable outcome of rational strategic behavior in a market that clears through a sequence of sales rather than through a Walrasian auctioneer. There are N potential buyers, each of whom is interested in securing one of K available units. Without fully modeling the selling side of the market, we assume it generates a single asking price at each instant of time according to some given function of buyer behavior to date. Each potential buyer observes all prices and all past offers to trade, and can accept the current asking price at any instant, in which case, supply permitting, the buyer trades at that price.

So traders have to decide both whether *and when* to offer to buy, all the while conditioning their strategies on the information that has been revealed in the market to date. Regarding the function generating the asking prices, we specify only that (i) if there is no demand at a price, then the next asking price is lower, and (ii) if demand exceeds remaining supply at any instant, then no trade actually takes place at that time but the next asking price is higher and only those who attempted to trade are allowed to buy subsequently.²² Note, however, that even if we did restrict attention to a specific price-setting process, the direct approach of computing buyers' optimal behavior using first-order conditions as a function of all prior behavior to solve a dynamic program would generally be completely intractable.

To use the RET we must first ensure that the appropriate assumptions are satisfied. We assume, of course, that buyers' valuations are independently drawn from a common, strictly increasing, atomless distribution, and that there is no discounting during the time the mechanism takes. And the objects do eventually go to the highest-valuation buyers, and the lowest-possible-valuation buyer makes zero surplus in equilibrium, because of our assumption that if demand ever exceeds remaining supply then no trade takes place and non-demanders are henceforth excluded. So

²¹ The same approach is also an economical method of computing equilibrium bids in many standard auctions. For example, in an ascending auction for a single unit, the expected payment of a bidder equals her probability of winning times the expected second-highest valuation among all the bidders conditional on her value being higher. So the RET implies that her equilibrium bid in a standard all-pay auction equals this quantity. Similarly, the RET implies that her equilibrium bid in a first-price sealed-bid auction equals the expected second-highest valuation among all the bidders conditional on her value being higher. See Appendix 1.A for more details and discussion.

²² Additional technical assumptions are required to ensure that all units are sold in finite time. See Bulow and Klemperer (1994) for full details.

the RET applies, and it also applies to any subgame of the whole game.²³

Under our assumptions, then, starting from any point of the process, the remainder of the game is revenue equivalent to what would result if the game were halted at that point and the remaining k objects were sold to the remaining buyers using a standard ascending auction (which sells all k objects at the $(k + 1)$ th highest valuation among the remaining buyers). So at any point of our game we know the expected payment of any buyer in the remainder of our game, and therefore also the buyer's expected payment conditional on winning.²⁴ But any potential buyer whose expected payment conditional on winning equals or exceeds the current asking price will attempt to buy at the current price.²⁵ This allows us to completely characterize buyer behavior, so fully characterizes the price path for any given rule generating the asking prices.

It is now straightforward to show (see Bulow and Klemperer, 1994) that potential buyers are extremely sensitive to the new information that the price process reveals. So almost any seller behavior—for example, starting at a very high price and slowly lowering the price continuously until all the units are sold or there is excess demand—will result in “frenzies” of trading activity in which many buyers bid simultaneously, even though there is zero probability that two buyers have the same valuation.²⁶ Furthermore these frenzies will

²³ If, instead, excess demand resulted in random rationing, the highest-valuation buyers might not win, violating the requirements of the RET, so even if we thought this was more natural it would make sense to begin with our assumption to be able to analyze and understand the process using the RET. The effects of the alternative assumption could then be analyzed with the benefit of the intuitions developed using the RET. Bulow and Klemperer (1994) proceed in exactly this way.

²⁴ Specifically, if k objects remain, the buyer's expected payment conditional on winning will be the expected $(k + 1)$ th highest valuation remaining conditional on the buyer having a valuation among the k highest remaining, and conditional on all the information revealed to date. This is exactly the buyer's expected payment conditional on winning an object in the ascending auction, since in both cases only winners pay and the probability of a bidder winning is the same.

²⁵ The marginal potential buyer, who is just indifferent about bidding now, will either win now or will never win an object. (If bidding now results in excess demand, this bidder will lose to inframarginal current bidders, since there is probability zero that two bidders have the same valuation.) So conditional on winning, this bidder's actual payment is the current price. Inframarginal bidders, whose expected payment conditional on winning exceeds the current price, may eventually end up winning an object at above the current price.

²⁶ To see why a frenzy must arise if the price is lowered continuously, note that for it to be rational for any potential buyer to jump in and bid first, there must be positive probability that there will be a frenzy large enough to create excess demand immediately following the first bid. Otherwise the strategy of waiting to bid until another player has bid first would guarantee a lower price.

For more general seller behavior, the point is that while buyers' valuations may be very dispersed, higher-valuation buyers are all almost certainly inframarginal in terms of whether to buy and are therefore all solving virtually identical optimization problems of when to buy. So a small change in asking price, or a small change in market conditions (such as the information revealed by a single trade) at a given price, can make a large number of buyers change from being unwilling to trade to wanting to trade.

The only selling process that can surely avoid a frenzy is a repeated Dutch auction.

sometimes lead to “crashes” in which it becomes common knowledge that the market price must fall a substantial distance before any further trade will take place.²⁷ Bulow and Klemperer also show that natural extensions to the model (e.g., “common values”, the possibility of resale, or an elastic supply of units) tend to accentuate frenzies and crashes. Frenzies and crashes arise precisely because buyers are rational and strategic; by contrast buyer irrationality might lead to “smoother” market behavior.

Of course our main point here is not the details of the process, but rather that the RET permits the solution and analysis of the dynamic price path of a market that would otherwise seem completely intractable to solve for.

2.3 TRANSLATING LOOSER ANALOGIES FROM AUCTIONS INTO ECONOMICS: ASCENDING VS. (FIRST-PRICE) SEALED-BID AUCTIONS

A major focus of auction theory has been contrasting the revenue and efficiency properties of “ascending” and “sealed-bid” auctions.²⁸ Ideas and intuitions developed in these comparisons have wide applicability.

2.3.1 *Internet Sales Versus Dealer Sales*

There is massive interest in the implications of e-commerce and internet sales. For example, the advent of internet sales in the automobile industry as a partial replacement for traditional methods of selling through dealers has been widely welcomed in Europe;²⁹ the organization of the European automobile market is currently a major policy concern both in official circles and the popular press, and the internet sales are seen as increasing “transparency”. But is transparency a good thing?

Auction theory shows that internet sales need *not* be good for consumers. Clearly transparent prices benefit consumers if they reduce consumers’ search costs so that in effect there are more competitors for every consumer.³⁰ And of course internet sales may also lower prices by cutting out the fixed costs of dealerships, albeit by also cutting out the additional services that dealers

²⁷ The price process is also extremely sensitive to buyer valuations; an arbitrarily small change in one buyer’s value can discontinuously and substantially change all subsequent trading prices.

²⁸ By “sealed-bid” we mean standard first-price sealed-bid auctions. “Ascending” auctions have similar properties to second-price sealed-bid auctions. See sections 1.1.2–1.1.4 for an introduction to the different types of auctions.

²⁹ See, for example, “May the net be with you”, *Financial Times*, 21 October 1999, p. 22. In the United Kingdom, Vauxhall began selling a limited number of special models over the internet late in 1999, while Ford began a pilot project in Finland.

³⁰ There may be both a direct effect (that consumers can observe more firms), and an indirect effect (that new entry is facilitated). See Baye and Morgan (2001) and Kühn and Vives (1995) for more discussion.

provide. But transparency also makes internet sales more like ascending auctions, by contrast with dealer sales that are more like (first-price) sealed-bid auctions, and we will show this is probably *bad* for consumers.

Transparent internet prices are readily observable by a firm's competitors so result, in effect, in an "ascending" auction; a firm knows if and when its offers are being beaten and can rapidly respond to its competitors' offers if it wishes. So, viewing each car sale as a separate auction, the price any consumer faces falls until all but one firm quits bidding to sell to him. (The price is, of course, descending because firms are competing to sell, but the process corresponds exactly to the standard ascending auction among bidders competing to buy an object, and we therefore maintain the standard "ascending" terminology.)

On the other hand, shopping to buy a car from one of competing dealers is very like procuring in a (first-price) "sealed-bid" auction. It is typically impossible to credibly communicate one dealer's offer to another. (Car dealers often deliberately make this hard by refusing to put an offer in writing.) So from the buyer's perspective it is as if sellers were independently making sealed-bid offers in ignorance of the competition.

Of course, the analogies are imperfect,³¹ but they serve as a starting point for analysis. So what does auction theory suggest?

Since, under the conditions of the revenue equivalence theorem, there is no difference between the auction forms for either consumer or producer welfare, we consider the implications of the most important violations of the conditions.

First, market demand is downward sloping, not inelastic.³² Hansen (1988) showed that this means consumers always prefer the sealed-bid setting, and firms may prefer it also; the sum of producer and consumer surpluses is always higher in a sealed-bid auction.³³ The intuition is that in an "ascending" auction the sales price equals the runner-up's cost, so is less reflective of the winner's cost than is the sealed-bid price. So the sealed-bid auction is more productively efficient (the quantity traded better reflects the winner's cost) and provides greater incentive for aggressive bidding (a more aggressive sealed bid not only increases the probability of winning, but also increases the quantity traded contingent on winning).

³¹ The analogies are less good for many other products. For lower-value products than cars, internet sales are less like an "ascending" auction since search costs will allow price dispersion, while traditional sales through posted prices in high-street stores are more like "ascending" auctions than are dealer sales of cars.

Note also that the outcomes of the two auction types differ most when competitors have private information about their costs, which is more likely when competitors are original manufacturers than when competitors are retailers selling goods bought at identical prices from the same wholesaler.

³² For an individual consumer, demand might be inelastic for a single car up to a reservation price. From the point of view of the sellers who do not know the consumer's reservation price, the expected market demand is downward sloping.

³³ Of course, Hansen is maintaining the other important assumptions of the revenue equivalence theorem.

Second, we need to consider the possibilities for collusion, implicit or explicit. The general conclusion is that ascending auctions are more susceptible to collusion, and this is particularly the case when, as in our example, many auctions of different car models and different consumers are taking place simultaneously.³⁴ As has been observed in the US and German auctions of radio spectrum, for example, bidders may be able to tacitly coordinate on dividing up the spoils in a simultaneous ascending auction. Bidders can use the early rounds when prices are still low³⁵ to signal their views about who should win which objects, and then, when consensus has been reached, tacitly agree to stop pushing prices up.³⁶ The same coordination cannot readily be achieved in simultaneous sealed-bid auctions, where there is neither the opportunity to signal, nor the ability to retaliate against a bidder who fails to cooperate.³⁷ The conclusion is less stark when there are many repetitions over time, but it probably remains true that coordination is easier in ascending auctions. Furthermore, as is already well understood in the industrial-organization literature,³⁸ this conclusion is strengthened by the different observabilities of internet and dealer sale prices which make mutual understanding of firms' strategies, including defections from "agreements", far greater in the internet case. So selling over the internet probably makes it easier for firms to collude.

A third important issue is that bidders may be asymmetric. Then "ascending" auctions are generally more efficient (because the lowest-cost bidders win³⁹), but

³⁴ See Robinson (1985) and Milgrom (1987) for discussion of the single-unit case, Menezes (1996), Weber (1997), Engelbrecht-Wiggans and Kahn (1998c), Ausubel and Schwartz (1999), Brusco and Lopomo (2002a), and Cramton and Schwartz (2000) for the multi-unit case. Sections 3.2 and 4.3.2 review these arguments and gives many examples.

³⁵ Bidders are competing to buy rather than sell spectrum, so prices are ascending rather than descending.

³⁶ For example, in a 1999 German spectrum auction Mannesman bid a low price for half the licenses and a slightly lower price for the other half. One of T-Mobile's managers said. "There were no agreements with Mannesman. But [T-Mobile] interpreted Mannesman's first bid as an offer." T-Mobile understood that it could raise the bid on the other half of the licenses slightly, and that the two companies would then "live and let live" with neither company challenging the other on "their" half. Just that happened. The auction closed after just two rounds with each of the bidders having half the licenses for the same low price (Jehiel and Moldovanu, 2000; Grimm, Riedel, and Wolfstetter, 2002).

In US FCC auctions, bidders have used the final three digits of multi-million dollar bids to signal the market ID codes of the areas they coveted, and a 1997 auction that was expected to raise \$1,800 million raised less than \$14 million. See Cramton and Schwartz (2001), and "Learning to Play the Game", *The Economist*, 17 May 1997, p. 120.

Sections 3.2 and 4.3.2 and chapter 5 give many more examples.

³⁷ The low prices in the ascending auction are supported by the threat that if a bidder overbids a competitor anywhere, then the competitor will retaliate by overbidding the first bidder on markets where the first bidder has the high bids.

³⁸ At least since Stigler (1964).

³⁹ To the extent that the auctions for individual consumers are independent single-unit auctions, an ascending auction is efficient under a broad class of assumptions if bidders' private signals are single-dimensional, even with asymmetries among bidders and common-value components to valuations. See Maskin (1992).

sealed-bid auctions typically yield lower consumer prices.⁴⁰ In this case economists generally favor ascending auctions, but competition-policy practitioners should usually prefer sealed-bid auctions because most competition regimes concentrate on consumer welfare.

Furthermore, this analysis ignores the impact of auction type on new entry in the presence of asymmetries. Because an “ascending” auction is generally efficient, a potential competitor with even a slightly higher cost (or lower quality) than an incumbent will see no point in entering the auction. However, the same competitor might enter a sealed-bid auction which gives a weaker bidder a shot at winning. The extra competition may lower prices very substantially. Of course the entry of the weaker competitor may also slightly reduce efficiency, but if competition is desirable per se, or if competition itself improves efficiency, or if the objective is consumer welfare rather than efficiency, then the case for sealed-bid auctions is very strong (see sections 2.3.2, 3.3, 3.5.2, and 4.3.1).

Although there are other dimensions in which our setting fails the revenue equivalence assumptions, they seem less important.⁴¹ So the transparency induced between firms that makes internet sales more like ascending auctions than sealed-bid auctions is probably bad for consumers. While gains from lower consumer search costs and dealer costs could certainly reverse this conclusion, auction-theoretic considerations mount a strong case against “transparent” internet sales.⁴²

In another application of auction-theoretic insights to e-commerce, Bulow and Klemperer (forthcoming) apply Milgrom and Weber’s (1982a) celebrated *linkage principle* to show when the price discrimination that internet markets make possible *helps* consumers.

⁴⁰ A price-minimizing auction allocates the object to the bidder with the lowest “virtual cost”, rather than to the one with the lowest actual cost. (See section 2.4; virtual cost is the analogous concept to marginal revenue for an auction to buy an object.) Compared to an ascending auction, a sealed-bid auction discriminates in favor of selling to “weaker” bidders, whose costs are drawn from higher distributions, because they bid more aggressively (closer to their actual costs) than stronger ones. But, for a given cost, a weaker bidder has a lower virtual cost than a stronger one. So in the sealed-bid auction often, but not always, yields lower prices. See section 1.7.1.

⁴¹ Other violations of the revenue equivalence assumptions may include buyer and seller risk aversion which both favor sealed-bid auctions, and affiliation of costs which favors ascending auctions.

⁴² Empirical evidence is limited. Lee (1998) and Lee, Westland, and Hong (1999) find electronic markets yield higher prices than conventional markets for cars. Scott Morton, Zettelmeyer, and Silva Riso (2001) find that California customers get lower prices if they use automobile internet sites, but this is unsurprising since these sites merely refer customers to dealers for price quotes, so behave more like traditional dealers than like the “transparent” sites that we have described and that are being promised in Europe.

2.3.2 *Anglo-Dutch Auctions, a Theory of Rationing, and Patent Races*

The last disadvantage of ascending auctions discussed above—the dampening effect on entry—has been very important in practical auction contexts (see parts C and D). For example, in the main (1995) auction of US mobile-phone licenses some large potential bidders such as MCI, the United States’s third-largest phone company, failed to enter at all, and many other bidders were deterred from competing seriously for particular licenses such as the Los Angeles and New York licenses which therefore sold at very low prices.⁴³ Entry was therefore a prominent concern when the United Kingdom planned an auction of four UMTS “third generation” mobile-phone licenses in 1998 for a market in which four companies operated mobile telephone services and therefore had clear advantages over any new entrant.⁴⁴

In this case the design chosen was an “Anglo-Dutch” auction as first proposed in Klemperer (1998):⁴⁵ in an Anglo-Dutch auction for four licenses the price rises continuously until five bidders remain (the “English” stage), after which the five survivors make sealed-bids (required to be no lower than the current price level) and the four winners pay the fourth-highest bid (the “Dutch” stage). Weak bidders have an incentive to enter such an auction because they know they have a chance of winning at the sealed-bid stage if they can survive to be among the five finalists. The design accepts some risk of an ex-post inefficient allocation in order to increase the chance of attracting the additional bidders that are necessary for a successful auction and reasonable revenues.^{46,47}

⁴³ See Klemperer and Pagnozzi (forthcoming) for econometric evidence of these kinds of problems in US spectrum auctions, Klemperer (1998) and Bulow and Klemperer (2002) for extensive discussion, and Bulow, Huang, and Klemperer (1999) for related modeling.

⁴⁴ Bidders could not be allowed to win more than one license each.

⁴⁵ See Klemperer (1998), Radiocommunications Agency (1998a,b) and sections 3.5.3 and 6.5.1 for more details and for variants on the basic design. (The Agency was advised by Binmore, Klemperer and others.)

⁴⁶ The additional bidders might yield a higher price even after the English stage, let alone after the final stage, than in a pure ascending auction.

⁴⁷ The design performed very successfully in laboratory testing, but the auction was delayed until 2000 and technological advances made it possible to offer five licenses, albeit of different sizes. The additional license resolved the problem of attracting new entrants, and since collusion was not a serious problem in this case (bidders were not allowed to win more than one license each), it was decided to switch to a simultaneous ascending design.

The actual UK auction was very successful, but the wisdom of the UK decision not to run an ascending auction when the number of strong bidders equaled the number of licenses was confirmed when the Netherlands did just this three months later, and raised little more than one-quarter of the per capita revenue raised by the United Kingdom. In large part the Netherlands’ problem was that their ascending auction deterred entry.

Denmark also had the same number of strong bidders as licenses, and (successfully) used a sealed-bid auction for similar reasons that the United Kingdom would have run an Anglo-Dutch auction in this context. (In Denmark it was clear that there were too few potential bidders to make an Anglo stage worthwhile.)

See parts C and D for more detail.

Translating this idea into a more traditional economics context suggests a theory of why firms might ration their output at prices at which there is excess demand as, for example, microprocessor manufacturers routinely do after the introduction of a new chip. Raising the price to clear the market would correspond to running an ascending auction. It would be ex-post efficient and ex-post profit maximizing, but would give poor incentives for weaker potential customers who fear being priced out of the market to make the investments necessary to enter the market (such as the product design necessary to use the new chip). Committing to rationing at a fixed price at which demand exceeds supply is ex-post inefficient,⁴⁸ but may encourage more entry into the market and so improve ex-ante profits. Details and more examples are in Gilbert and Klemperer (2000).

A similar point is that a weaker firm may not be willing to enter a patent race in which all parties can observe others' progress. Such a race is akin to an ascending auction in which a stronger rival can always observe and overtake a weaker firm which therefore has no chance of winning.⁴⁹ A race in which rivals' progress cannot be monitored is more akin to a sealed-bid auction and may attract more entry.

These analogies illustrate how an insight that is routine in auction theory may help develop ideas in economics more broadly.

2.4 EXPLOITING DEEPER CONNECTIONS BETWEEN AUCTIONS AND ECONOMICS: MARGINAL REVENUES

The previous sections showed how a variety of economic problems can be thought of in auction-theoretic terms, allowing us to use tools such as the revenue equivalence theorem and intuitions such as those from the comparison of ascending and sealed-bid auctions. This section explains that the connections between auction theory and standard economic theory run much deeper.

Much of the analysis of optimal auctions can be phrased, like the analysis of monopoly, in terms of "marginal revenues". Imagine a firm whose demand curve is constructed from an arbitrarily large number of bidders whose values are independently drawn from a bidder's value distribution. When bidders have independent private values, a bidder's "marginal revenue" is defined as the marginal revenue of this firm at the price that equals the bidder's actual value (see figure 2.1).⁵⁰

⁴⁸ We assume any resale is inefficient. But see Cramton, Gibbons, and Klemperer (1987).

⁴⁹ Of course, this point is closely related to the idea of " ϵ -preemption" in R&D races with observability that has already been well discussed in the standard industrial organization literature (Fudenberg, Gilbert, Stiglitz, and Tirole, 1983).

⁵⁰ The point of this construction is particularly clear when a seller faces a single bidder whose private value is distributed according to $F(v)$. Then setting a take-it-or-leave-it price of v yields expected sales, or "demand", $1 - F(v)$, expected revenue of $v(1 - F(v))$ and expected marginal revenue $d(qv)/dq = v - [1 - F(v)]/f(v)$. See Appendix 1.B.

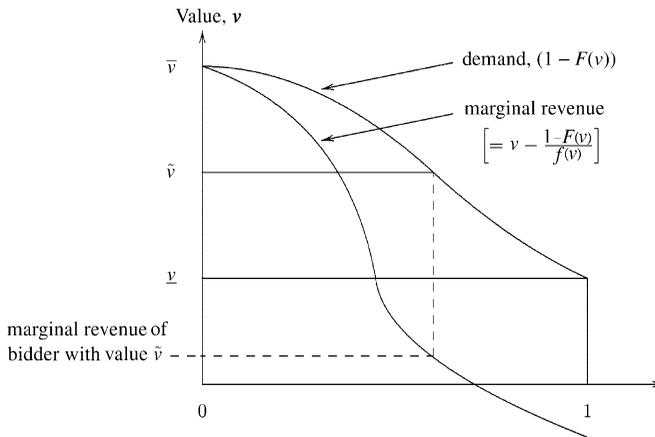


Figure 2.1 Construction of marginal revenue of bidder with value \tilde{v} drawn from distribution $F(v)$ on $[\underline{v}, \bar{v}]$.

Although it had been hinted at before,⁵¹ the key point was first explicitly drawn out by Bulow and Roberts (1989) who showed that under the assumptions of the revenue equivalence theorem *the expected revenue from an auction equals the expected marginal revenue of the winning bidder(s)*. The new results in the article were few—the paper largely mimicked Myerson (1981) while renaming Myerson’s concept of “virtual utility” as “marginal revenue”^{52,53}—but their contribution was nevertheless important. Once the connection had been made it was possible to take ways of thinking that are second nature to economists from the standard theory of monopoly pricing and apply them to auction theory.

For example, once the basic result above (that an auction’s expected revenue equals the winning bidder’s expected marginal revenue) was seen,

⁵¹ For example, Mussa and Rosen’s (1978) analysis of monopoly and product quality contained expressions for ‘marginal revenue’ that look like Myerson’s (1981) analysis of optimal auctions.

⁵² Myerson’s results initially seemed unfamiliar to economists in part because his basic analysis (although not all his expressions) expressed virtual utilities as a function of bidders’ values, which correspond to *prices*, and so computed revenues by integrating along the vertical axis, whereas we usually solve monopoly problems by expressing marginal revenues as functions of *quantities* and integrating along the horizontal axis of the standard (for monopoly) picture.

⁵³ Bulow and Roberts emphasize the close parallel between a monopolist third-degree price-discriminating across markets with different demand curves, and an auctioneer selling to bidders whose valuations are drawn from different distributions. For the {monopolist / auctioneer}, {expected revenue} is maximized by selling to the {consumers / bidder} with the highest marginal revenue(s), *not* necessarily the highest value(s), subject to never selling to a {consumer / bidder} with marginal revenue less than the {monopolist’s marginal cost / auctioneer’s own valuation}, assuming (i) resale can be prohibited, (ii) credible commitment can be made to {sticking to any reserve price / no future sales}, and (iii) {higher “types” of any bidder have higher marginal revenues than lower “types” of the same bidder / marginal revenue curves are all downward sloping}, etc.

Bulow and Klemperer (1996) were able to use a simple monopoly diagram to derive it more simply and under a broader class of assumptions than had previously been done by Myerson or Bulow and Roberts.⁵⁴ Bulow and Klemperer also used standard monopoly intuition to derive additional results in auction theory.

The main benefits from the marginal-revenue connection come from translating ideas from monopoly analysis into auction analysis, since most economists' intuition for and understanding of monopoly is much more highly developed than for auctions. But it is possible to go in the other direction too, from auction theory to monopoly theory.

Consider, for example, the main result of Bulow and Klemperer (1996).

Proposition (Auction-Theoretic Version). *An optimal auction of K units to Q bidders earns less profit than a simple ascending auction (without a reserve price) of K units to $Q + K$ bidders, assuming (a) bidders are symmetric, (b) bidders are serious (i.e., their lowest-possible valuations exceed the seller's supply cost), and (c) bidders with higher valuations have higher marginal revenues.*⁵⁵

Proof. See Bulow and Klemperer (1996).

Application. One application is to selling a firm (so $K = 1$). Since the seller can always resort to an ascending auction, attracting a single additional bidder is worth more than any amount of negotiating skill or bargaining power against an existing bidder or bidders, under reasonable assumptions. So there is little justification for, for example, accepting a "lock-up" bid for a company without fully exploring the interest of alternative possible purchasers.

The optimal auction translates, for large Q and K , to the monopolist's optimum. An ascending auction translates to the competitive outcome, in which price-taking firms make positive profits only because of the fixed supply of units. (An ascending auction yields the $(K + 1)$ th highest value among the bidders; in a perfectly competitive market an inelastic supply of K units is in equilibrium with demand at any price between the K th and $(K + 1)$ th highest value, but the distinction is unimportant for large K .) So one way of expressing the result in the market context is:

⁵⁴ See Appendix 1.B for an exposition.

⁵⁵ See Bulow and Klemperer (1996) for a precise statement. We do not require bidders' valuations to be private, but do place some restrictions on the class of possible mechanisms from which the "optimal" one is selected, if bidders are not risk-neutral or their signals are not independent. We assume bidders demand a single unit each.

Proposition (Monopoly-Theoretic Version). *A perfectly competitive industry with (fixed) capacity K and Q consumers would gain less by fully cartelizing the industry (and charging the monopoly price) than it would gain by attracting K new potential customers into the industry with no change in the intensity of competition, assuming (a') the K new potential consumers have the same distribution of valuations as the existing consumers, (b') all consumers' valuations for the product exceed sellers' supply costs (up to sellers' capacity), and (c') the marginal–revenue curve constructed from the market–demand curve is downward sloping.⁵⁶*

Proof. No proof is required—the proposition is implied by the auction-theoretic version—but once we know the result we are looking for and the necessary assumptions, it is very simple to prove it directly using introductory undergraduate economics and we do this in a brief Appendix 2.B.

Application. One application is that this provides conditions under which a joint-marketing agency does better to focus on actually marketing rather than (as some of the industrial organization literature suggests) on facilitating collusive practices.⁵⁷

2.5 APPLYING AUCTION THEORY TO PRICE-SETTING OLIGOPOLIES

We have stressed the applications of auction theory to contexts that might not be thought of as auctions, but even though price-setting oligopolies are obviously auctions, the insights that can be obtained by thinking of them in this way are often passed by.

2.5.1 Marginal-Cost Pricing is NOT the Unique Bertrand Equilibrium

One of the most famous results in economics is the “Bertrand paradox” that with just two firms with constant and equal marginal costs in a homogeneous-products industry the unique equilibrium is for both firms to set price equal to marginal cost and firms earn zero profit. This “theorem” is widely quoted in standard texts. But it is *false*. There are other equilibria with large profits, for some standard demand curves, a fact that seems until recently to have been known only to a few auction theorists.⁵⁸

⁵⁶ We are measuring capacity in units such that each consumer demands a single unit of output. Appendix 2.B makes it clear how the result generalizes.

⁵⁷ Of course the agency may wish to pursue both strategies in practice.

⁵⁸ We assume firms can choose any prices. It is well known that if prices can only be quoted in whole pennies, there is an equilibrium with positive (but small) profits in which each firm charges one penny above cost. (With perfectly inelastic demand, there is also an equilibrium in which each firm charges two pennies above cost.)

Auction theorists are familiar with the fact that a boundary condition is necessary to solve a sealed-bid auction. Usually this is imposed by assuming no bidder can bid less than any bidder's lowest-possible valuation, but there are generally a continuum of equilibria if arbitrarily negative bids are permitted.⁵⁹ Exactly conversely, with perfectly inelastic demand for one unit and, for example, two risk-neutral sellers with zero costs, it is a mixed-strategy equilibrium for each firm to bid above any price p with probability k/p , for any fixed k . (Each firm therefore faces expected residual demand of constant elasticity -1 , and is therefore indifferent about mixing in this way; profits are k per firm.)

It is not hard to see that a similar construction is possible with downward-sloping demand, for example, standard constant-elasticity demand, provided that monopoly profits are unbounded. (See especially, Baye and Morgan (1999a) and Kaplan and Wettstein (2000)). One point of view is that the non-uniqueness of the "Bertrand paradox" equilibrium is a merely technical point since it requires "unreasonable" (even though often assumed⁶⁰) demand. However, the construction immediately suggests another more important result: quite generally (including for demand which becomes zero at some finite choke price) there are very profitable mixed-strategy ε -equilibria to the Bertrand game, even though there are no pure-strategy ε -equilibria. That is, there are mixed strategies that are very different from marginal-cost pricing in which no player can gain more than a very small amount, ε , by deviating from the strategies.⁶¹ (There are also "quantal response" equilibria with a similar flavor.) Experimental evidence suggests that these strategies may be empirically relevant (see Baye and Morgan, 1999b).⁶²

2.5.2 *The Value of New Consumers*

The Revenue Equivalence Theorem (RET) can of course be applied to price-setting oligopolies.⁶³

⁵⁹ For example, if each of two risk-neutral bidders' private values is independently drawn from a uniform distribution on the open interval $(0, 1)$ then for any non-negative k there is an equilibrium in which a player with value v bids $(v/2) - (k/v)$. If it is common knowledge that both bidders have value zero, there is an equilibrium in which each player bids below any price $-p$ with probability k/p , for any fixed non-negative k .

⁶⁰ This demand can, for example, yield unique and finite-profit Cournot equilibrium.

⁶¹ Of course, the concept of mixed-strategy ε equilibrium used here is even more contentious than either mixed-strategy (Nash) equilibria or (pure-strategy) ε equilibrium. The best defense for it may be its practical usefulness.

⁶² Spulber (1995) uses the analogy with a sealed-bid auction to analyze a price-setting oligopoly in which, by contrast with our discussion, firms do not know their rivals' costs. For a related application of auction theory to price-setting oligopoly, see Athey, Bagwell, and Sanchirico (forthcoming).

⁶³ See section 2.2 for discussion of the RET. As another example of its application to price-setting oligopolies, Vives (2002) uses the revenue equivalence theorem to compare oligopoly equilibria with incomplete and complete (or shared) information about firms' constant marginal costs, and so shows information sharing is socially undesirable in this context.

For example: what is the value of new consumers in a market with strong brand loyalty? If firms can price discriminate between new uncommitted consumers and old “locked-in” consumers, Bertrand competition for the former will mean their value is low, but what if price discrimination is impossible?

In particular, it is often argued that new youth smokers are very valuable to the tobacco industry because brand loyalty (as well as loyalty to the product) is very high (only about 10 percent of smokers switch brands in any year), so price-cost margins on all consumers are very high. Is there any truth to this view?

The answer, of course, under appropriate assumptions, is that the RET implies that the ability to price discriminate is irrelevant to the value of the new consumers: With price discrimination, we can model the oligopolists as acting as monopolists against their old customers, and as being in an “ascending”⁶⁴ price auction for the uncommitted consumers with the firm which is prepared to price the lowest selling to all these consumers at the cost of the runner-up firm. Alternatively, we can model the oligopolists as making sealed bids for the uncommitted consumers with the lowest bidder selling to these consumers at its asking price. The expected profits are the same under the RET assumptions. Absent price discrimination, a natural model is the latter one, but in addition each oligopolist must discount its price to its own locked-in customers down to the price it bids for the uncommitted consumers. The RET tells us that the total cost to the industry of these “discounts” to old consumers will on average precisely compensate the higher sale price achieved on new consumers.⁶⁵ That is, the net value to the industry of the new consumers is exactly as if there was Bertrand competition for them, even when the inability to price discriminate prevents this.

So Bulow and Klemperer (1998) argue that the economic importance to the tobacco companies of the youth market is actually very tiny, even

⁶⁴ The price is descending because the oligopolists are competing to sell rather than buy, but it corresponds to an ascending auction in which firms are competing to buy, and we stick with this terminology as in section 2.3.1.

⁶⁵ Specifically let n “old” consumers be attached to each firm i , and firms’ costs c_i be independently drawn from a common, strictly increasing, atomless distribution. There are m “new” consumers who will buy from the cheapest firm. All consumers have reservation price r .

Think of firms competing for the prize of selling to the new consumers, worth $m(r - c_i)$ to firm i . Firms set prices $p_i = r - d_i$ to “new” consumers; equivalently they set “discounts” d_i to consumers’ reservation prices. If price discrimination is feasible, the winner pays md_i for the prize and all firms sell to their old consumers at r . Absent price discrimination, the prices p_i apply to all firms’ sales, so relative to selling just to old consumers at price r , the winner pays $(m + n)d_i$ for the prize and the losers pay nd_i each.

For the usual reasons, the two sets of payment rules are revenue equivalent. For more discussion of this result, including its robustness to multi-period contexts, see Bulow and Klemperer (1998); if the total demand of new consumers is more elastic, their economic value will be somewhat less than our model suggests; for a fuller discussion of the effects of “brand loyalty” or “switching costs” in oligopoly see, especially, Klemperer (1987a,b, 1995) and Beggs and Klemperer (1992).

though from an accounting perspective new consumers appear as valuable as any others.⁶⁶

Similarly, applying the same logic to an international trade question, the value of a free-trading market to firms each of which has a protected home market is independent (under appropriate assumptions) of whether the firms can price discriminate between markets.⁶⁷

Section 2.3.1's discussion of oligopolistic e-competition develops this kind of analysis further by considering implications of failures of the RET.

2.5.3 Information Aggregation in Perfect Competition

Although the examples above, and in section 2.3,⁶⁸ suggest auction theory has been underused in analyzing oligopolistic competition, it has been very important in influencing economists' ideas about the limit as the number of firms becomes large.

An important strand of the auction literature has focused on the properties of pure-common-value auctions as the number of bidders becomes large, and asked: does the sale price converge to the true value, thus fully aggregating all of the economy's information even though each bidder has only partial information? Wilson (1977) and Milgrom (1979) showed assumptions under which the answer is "yes" for a first-price sealed-bid auction, and Milgrom (1981) obtained similar results for a second-price auction (or for a $(k + 1)$ th price auction for k objects).⁶⁹ So these models justify some of our ideas about perfect competition.

2.6 APPLYING AUCTION THEORY (AND ECONOMICS) TO AUCTION MARKETS

Finally, although it has not always been grasped by practitioners, some markets are literally auctions. The increasing recognition that many real markets are best understood through the lens of auction theory has stimulated a burst of new theorizing,⁷⁰ and created the new subject of market design that stands in similar relation to auction theory as engineering does to physics.

⁶⁶ If industry executives seem to value the youth segment, it is probably due more to concern for their own future jobs than concern for their shareholders.

⁶⁷ See also Rosenthal (1980).

⁶⁸ Bulow and Klemperer (forthcoming) provides an additional example.

⁶⁹ Matthews (1984), on the other hand, showed that the (first-price) sale price does not in general converge to the true value when each bidder can acquire information at a cost.

Pesendorfer and Swinkels (1997) recently breathed new life into this literature, by showing convergence under weaker assumptions than previously if the number of objects for sale, as well as the number of bidders, becomes large. See also Pesendorfer and Swinkels (2000), Swinkels (2001), and Kremer (2002).

⁷⁰ Especially on multi-unit auctions in which bidders are not restricted to winning a single unit each, since most markets are of this kind.

2.6.1 Important Auction Markets

It was not initially well understood that deregulated *electricity markets*, such as in the United Kingdom, are best described and analyzed as auctions of infinitely divisible quantities of homogeneous units.⁷¹ Although much of the early analysis of the UK market was based on Klemperer and Meyer (1989), which explicitly followed Wilson's (1979) seminal contribution to multi-unit auctions, the Klemperer and Meyer model was not thought of as an "auctions" paper and only recently received much attention among auction theorists.⁷² Indeed von der Fehr and Harbord (1993) were seen as rather novel in pointing out that the new electricity markets could be viewed as auctions. Now, however, it is uncontroversial that these markets are best understood through auction theory, and electricity market design has become the province of leading auction theorists, such as Wilson, who have been very influential.

Treasury bill auctions, like electricity markets, trade a divisible homogeneous good, but although treasury auctions have always been clearly understood to be "auctions", and the existing auction theory is probably even more relevant to treasury markets than to electricity markets,⁷³ auction theorists have never been as influential as they now are in energy markets. In part this is because the treasury auctions predated any relevant theory,⁷⁴ and the auctions seemed not to have serious problems. In part it may be because no clear view has emerged about the best form of auction to use; indeed one possibility is that the differences between the main types of auction may not be too important in this context (see section 3.4.6).⁷⁵

⁷¹ von der Fehr and Harbord (1998) provide a useful overview of electricity markets.

⁷² Klemperer and Meyer (1989) was couched as a traditional industrial organization study of whether competition is more like Bertrand or Cournot, following Klemperer and Meyer (1986).

⁷³ Non auction-theoretic issues which limit the direct application of auction theory to electricity markets include the very high frequency of repetition among market participants who have stable and predictable requirements which makes the theory of collusion in repeated games also very relevant; the nature of the game the major electricity suppliers are playing with the industry regulator who may step in and attempt to change the rules (again) if the companies are perceived to be making excessive profits; the conditions for new entry; and the effects of vertical integration of industry participants.

On the other hand, the interaction of a treasury auction with the financial markets for trading the bills both before and after the auction complicates the analysis of that auction.

⁷⁴ By contrast, the current UK government sales of gold are a new development, and government agencies have now consulted auction theorists (including myself) about the sale method.

⁷⁵ In a further interesting contrast the UK electricity market—the first major market in the world to be deregulated and run as an auction—was set up as a uniform price auction, but its perceived poor performance has led to a planned switch to an exchange market followed by a discriminatory auction (see sections 3.2 and 3.5.2; Office of Gas and Electricity Markets, 1999; Newbery, 1998; Wolfram, 1998, 1999). Meanwhile the vast majority of the world's treasury bill markets have until recently been run as discriminatory auctions (see Bartolini and Cottarelli, 1997), but the United States switched to uniform price auctions in late 1998 and several other countries have been experimenting with these. In fact, it seems unlikely that either form of auction is best either for all electricity markets or for all treasury markets (see, e.g., Klemperer, 1999b; Federico and Rahman, 2000; McAdams, 1998; Nyborg and Sundaresan, 1996).

Academics were involved at all stages of the *radio spectrum auctions* from suggesting the original designs to advising bidders on their strategies. The original US proponents of an auction format saw it as a complex environment that needed academic input, and a pattern of using academic consultants was set in the United States and spread to other countries (see part D).⁷⁶

Many other new auction markets are currently being created using the *Internet*, such as the online consumer auctions run by eBay, Amazon, and others which have over 10 million customers, and the business-to-business autoparts auctions being planned by General Motors, Ford, and Daimler-Chrysler which are expected to handle \$250 million in transactions a year. Here too auction theorists have been in heavy demand, and there is considerable ongoing experimentation with different auction forms.⁷⁷ Furthermore, we have already argued that internet markets that are not usually thought of as auctions can be illuminated by auction theory (see section 2.3.1 and Bulow and Klemperer, forthcoming).

2.6.2 Applying Economics to Auction Design

While many economic markets are now fruitfully analyzed as auctions, the most significant problems in auction markets and auction design are probably those with which industry regulators and competition authorities have traditionally been concerned—discouraging collusive, predatory and entry-detering behavior, and analyzing the merits of mergers or other changes to market structure.

⁷⁶ Evan Kwerel was especially important in promoting the use of auctions. The dominant design has been the simultaneous ascending auction sketched by Vickrey (1976), and proposed and developed by McAfee, Milgrom, and Wilson for the US auctions (see McMillan, 1994; McAfee and McMillan, 1996; and especially Milgrom, 2004). Although some problems have emerged, primarily its susceptibility to collusion and its inhospitability to entry (see section 2.3.2), it has generally been considered a success in most of its applications (see, e.g., Board, 1999; Cramton, 1997; Plott, 1997; Salant, 1997; Weber, 1997; Zheng, 2001).

A large part of the motivation for the US design was the possibility of complementarities between licenses (see Ausubel, Cramton, McAfee, and McMillan, 1997), although it is unproven either that the design was especially helpful in allowing bidders to aggregate efficient packages, or that it would work well if complementarities had been very significant. Ironically, the simultaneous ascending auction is most attractive when each of an exogenously fixed number of bidders has a privately known value for each of a collection of heterogenous objects, but (contrary to the US case) is restricted to buying at most a single license. In this case entry is not an issue, collusion is very unlikely, and the outcome is efficient. For this reason a version of the simultaneous ascending auction was designed by Binmore and Klemperer for the UK 3G auction (in which each bidder was restricted to a single license) after concerns about entry had been laid to rest.

A sealed-bid design was recently used very successfully in Denmark where attracting entry was a serious concern (see section 2.3.2). (See part D for discussion of the recent European spectrum auctions.)

⁷⁷ See, for example, Hall (2001). The UK government recently used the internet to run the world's first auction for greenhouse gas emissions reductions. (Peter Cramton, Eric Maskin and I advised on the design, and Larry Ausubel and Jeremy Bulow also helped with the implementation.)

This contrasts with most of the auction literature which focuses on Nash equilibria in one-shot games with a fixed number of bidders, and emphasizes issues such as the effects of risk aversion, correlation of information, budget constraints, complementarities, asymmetries, etc. While these are also important topics—and auction theorists have made important progress on them which other economic theory can learn from—they are probably not the main issues.

Although the relative thinness of the auction-theoretic literature on collusion and entry deterrence may be defensible to the extent general economic principles apply, there is a real danger that auction theorists will underemphasize these problems in applications. In particular, ascending, second-price, and uniform-price auction forms, while attractive in many auction theorists' models, are more vulnerable to collusive and predatory behavior than (first-price) sealed-bid and hybrid forms such as the Anglo-Dutch auction described in sections 2.3.2, 3.5.3, and 6.5.1. Part C provides an extensive discussion of these issues.

While auction theorists are justly proud of how much they can teach economics, they must not forget that the classical lessons of economics continue to apply.

2.7 CONCLUSION

Auction theory is a central part of economics, and should be a part of every economist's armory; auction theorists' ways of thinking shed light on a whole range of economic topics.

We have shown that many economic questions that do not at first sight seem related to auctions can be recast to be solvable using auction-theoretic techniques such as the revenue equivalence theorem. The close parallels between auction theory and standard price theory—such as those between the theories of optimal auctions and of price discrimination—mean ideas can be arbitrated from auction theory to standard economics, and vice versa. And the insights and intuitions that auction theorists have developed in comparing different auction forms can find fertile application in many other contexts.

Furthermore, while standard auction theory models already provide the basis of much work in labor economics, political economy, finance, and industrial organization, we have used the example of price-setting oligopoly to show that a much greater application of auction-theoretic thinking may be possible in these more obvious fields.

“Heineken refreshes the parts other beers cannot reach” was recently voted one of the top advertising campaigns of all time, worldwide. The moral of this chapter is that “Auction theory refreshes the parts other economics cannot reach.” Like Heineken, auction theory is a potent brew that we should all imbibe.

APPENDIX 2.A: COMPARING LITIGATION SYSTEMS

Assume that after transfers between the parties, the loser ends up paying fraction $\alpha \geq 0$ of his own expenses and fraction $\beta \leq 1$ of his opponent's. (The winner pays the remainder.)⁷⁸ So the American system is $\alpha = 1$, $\beta = 0$, the British system is $\alpha = \beta = 1$, the Netherlands system is roughly, $\alpha = 1$, $0 < \beta < 1$, and Quayle's is $\alpha = 2$, $\beta = 0$. It is also interesting to consider a "reverse-Quayle" rule $\alpha = 1$, $\beta < 0$ in which both parties pay their own expenses but the winner transfers an amount proportional to her own expenses to the loser. Let L be the average legal expenses spent per player.

The following slight generalization of the RET is the key: assuming the conditions of the RET all hold except for assumption (ii) (that is, the expected surplus of a bidder with the lowest-feasible valuation, say \underline{S} , may not be zero), it remains true that the expected surplus of any other types of bidder is a fixed amount above \underline{S} (see, e.g., Appendix 1.A); the fixed amount depends on the distribution of the parties' valuations, but unlike \underline{S} and L does not depend on the mechanism $\{\alpha, \beta\}$.

It follows that the average bidder surplus is \underline{S} plus a constant. But the average bidder surplus equals the average lawsuit winnings (expectation of $\{\text{probability of winning}\} \times \{\text{valuation}\}$) minus L , equals a constant minus L by assumption (i) of the RET. So $\underline{S} = K - L$ in which K is a constant independent of α and β . But since the lowest-valuation type always loses in equilibrium (by assumption (i) of the RET) she bids zero so $\underline{S} = -\beta L$ because in a one-shot game her opponent, on average, incurs expenses of L . Solving, $L = K/(1 - \beta)$ and the expected surplus of any given party is a constant minus $\beta K/(1 - \beta)$.

It follows that both expected total expenses and any party's expected payoff are invariant to α , hence the remarks in the text about the Quayle proposal. But legal expenses are increasing in β , indeed become unbounded in the limit corresponding to the British system. So the mechanism that minimizes legal expenses taking the set of lawsuits as given is the reverse-Quayle. The intuition is that it both increases the marginal cost of spending on a lawsuit and reduces the value of winning the suit. On the other hand, of course, bringing lawsuits becomes more attractive as β falls.

⁷⁸ As in the main text, we assume a symmetric equilibrium with strictly increasing bidding functions. For extreme values of α and β this may not exist (and we cannot then use the RET directly). See Baye, Kovenock, and de Vries (1997) for explicit solutions for the equilibria for different α and β .

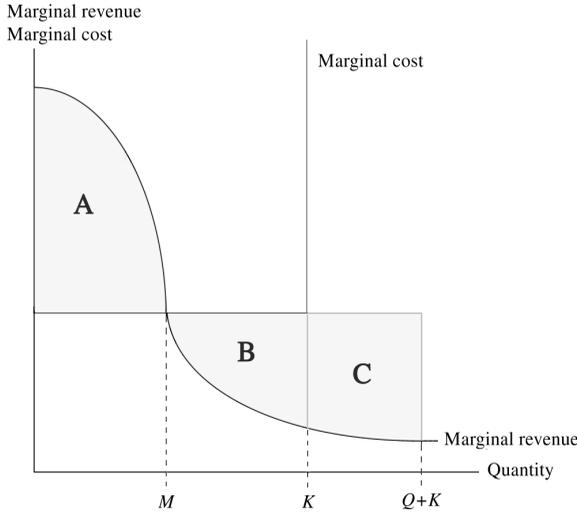


Figure 2.2 Marginal revenue if demand is expanded.

APPENDIX 2.B: DIRECT PROOF OF MONOPOLY-THEORETIC VERSION OF PROPOSITION IN SECTION 2.4

The proof rests precisely on the assumptions (a'), (b'), and (c'). Without loss of generality let firms' marginal costs be flat up to capacity,⁷⁹ and consider what would be the marginal revenue curve for the market if the K new consumers were attracted into it (see figure 2.2).

A monopolist on this (expanded) market would earn area A in profits, that is, the area between the marginal revenue and marginal cost curves up to the monopoly point, M . The perfectly competitive industry in the same (expanded) market would earn $\Pi^c = A - B$, that is, the integral of marginal revenue less marginal cost up to industry capacity, K . By assumption (a'), a monopolist (or fully cartelized industry) in the original market would earn $\Pi^m = [Q/(Q + K)]A$. Now the average marginal revenue up to quantity $Q + K$ equals the price at demand $Q + K$ (because total marginal revenue = price \times quantity), which exceeds marginal cost by assumption (b'), so $B + C \leq A$. Furthermore, by assumption (c'), and elementary geometry, $B \leq [(K - M)/(Q + K) - M](B + C)$. So $B \leq [(K - M)/(Q + K - M)]A$, and therefore $\Pi^c = A - B \geq Q/(Q + K - M)A \geq \Pi^m$, as required.

⁷⁹ If the industry cost curve is not flat up to the capacity, then use the argument in the text to prove the result for a cost curve that is flat and everywhere weakly above the actual cost curve. A fortiori, this proves the result for the actual curve, since a monopoly saves less from a lower cost curve than a competitive industry saves from the lower cost curve.