

PUTTING AUCTION THEORY TO WORK

PAUL MILGROM

Stanford University



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Paul Milgrom 2004

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2004

Printed in the United States of America

Typefaces Utopia 9.75/13.5 pt. and ITC Symbol *System* $\text{\LaTeX}2_{\epsilon}$ [TB]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Milgrom, Paul R. (Paul Robert), 1948–

Putting auction theory to work / Paul Milgrom.

p. cm. – (Churchill lectures in economics)

ISBN 0-521-55184-6 (hard) – ISBN 0-521-53672-3 (pbk.)

1. Auctions – Mathematical models. I. Title. II. Series.

HF5476.M55 2003

381'.17'01 – dc21

2003051544

ISBN 0 521 55184 6 hardback

ISBN 0 521 53672 3 paperback

Contents

<i>Preface</i>	<i>page xi</i>
<i>Foreword by Evan Kwerel</i>	xv
1 Getting to Work	1
1.1 Politics Sets the Stage	3
1.2 Designing for Multiple Goals	3
1.2.1 Substitutes and Complements	6
1.2.2 New Zealand's Rights Auction	9
1.2.3 Better Auction Designs	13
1.2.4 The FCC Design and Its Progeny	13
1.3 Comparing Seller Revenues	16
1.4 The Academic Critics	19
1.4.1 Resale and the Coase Theorem	19
1.4.2 Mechanism Design Theory	21
1.4.3 Theory and Experiment	25
1.4.4 Practical Concerns	26
1.5 Plan for This Book	31
PART I THE MECHANISM DESIGN APPROACH	35
2 Vickrey–Clarke–Groves Mechanisms	45
2.1 Formulation	45
2.2 Always Optimal and Weakly Dominant Strategies	49
2.3 Balancing the Budget	53
2.4 Uniqueness	55
2.5 Disadvantages of the Vickrey Auction	56
2.5.1 Practical Disadvantages	56
2.5.2 Monotonicity Problems	57
2.5.3 The Merger–Investment Disadvantage	60
2.6 Conclusion	61
	vii

3 The Envelope Theorem and Payoff Equivalence	64
3.1 Hotelling's Lemma	65
3.2 The Envelope Theorem in Integral Form	66
3.3 Quasi-linear Payoffs	69
3.3.1 Holmstrom's Lemma	70
3.3.2 The Green–Laffont–Holmstrom Theorem	71
3.3.3 Myerson's Lemma	73
3.3.4 Revenue Equivalence Theorems	75
3.3.5 The Myerson–Satterthwaite Theorem	77
3.3.6 The Jehiel–Moldovanu Impossibility Theorems	80
3.3.7 Myerson and Riley–Samuelson Revenue-Maximizing Auctions	84
3.3.8 The McAfee–McMillan Weak-Cartels Theorem	87
3.3.9 Sequential Auctions and Weber's Martingale Theorem	90
3.3.10 Matthews Theorem: Risk Averse Payoff Equivalence	91
3.4 Conclusion	94
4 Bidding Equilibrium and Revenue Differences	98
4.1 The Single Crossing Conditions	99
4.1.1 The Monotonic Selection Theorem	101
4.1.2 The Sufficiency Theorem	102
4.1.3 The Constraint Simplification Theorem	105
4.1.4 The Mirrlees–Spence Representation Theorem	106
4.2 Deriving and Verifying Equilibrium Strategies	110
4.2.1 The Second-Price Auction with a Reserve Price	111
4.2.2 The Sealed Tender, or First-Price, Auction	112
4.2.3 The War of Attrition Auction	117
4.2.4 The All-Pay Auction	119
4.3 Revenue Comparisons in the Benchmark Model	119
4.3.1 Payoff Equivalence without Revenue Equivalence	121
4.3.2 Budget Constraints	132
4.3.3 Endogenous Quantities	135
4.3.4 Correlated Types	137
4.4 Expected-Revenue-Maximizing Auctions	140
4.4.1 Myerson's Theorem	144
4.4.2 Bulow–Klemperer Theorem	148
4.4.3 The Irregular Case	148
4.5 Auctions with Weak and Strong Bidders	149
4.6 Conclusion	154

5 Interdependence of Types and Values	157
5.1 Which Models and Assumptions are “Useful”?	158
5.1.1 Payoffs Depend Only on Bids and Types	158
5.1.2 Types Are One-Dimensional and Values Are Private	159
5.1.3 Types Are Statistically Independent	161
5.2 Statistical Dependence and Revenue-Maximizing Auctions	162
5.3 Wilson’s Drainage Tract Model	166
5.3.1 Equilibrium	167
5.3.2 Profits and Revenues	173
5.3.3 Bidder Information Policy	175
5.3.4 Seller Information Policy	177
5.4 Correlated Types and Interdependent Values	181
5.4.1 Affiliation	182
5.4.2 The Milgrom–Weber Ascending Auction Models	187
5.4.2.1 The (Second-Price) Button Auction with Minimal Information	188
5.4.2.2 The Button Auction with Maximal Information	195
5.4.2.3 Some Revenue Comparisons	198
5.4.3 First-Price Auctions	200
5.5 Conclusion	204
6 Auctions in Context	208
6.1 The Profit and Surplus Contribution of an Entrant	214
6.2 Symmetric Models with Costly Entry	216
6.2.1 Symmetric Bidders and Uncoordinated Entry	218
6.2.1.1 Equilibrium in Entry and Bidding Decisions	218
6.2.1.2 Setting the Reserve Price	222
6.2.2 Coordinating Entry among Symmetric Competitors	225
6.2.2.1 Pre-qualifying Bidders	227
6.2.2.2 Auctions, Negotiations, and Posted Prices	230
6.2.2.3 Buy Prices	232
6.3 Asymmetric Models: Devices to Promote Competition	234
6.3.1 Example of Set-asides	235
6.3.2 Example of Bidding Credits	237
6.3.3 Example of Lot Structure and Consolation Prizes	238
6.3.4 Premium Auctions	239
6.3.5 Dutch vs. English Auctions and the Anglo-Dutch Design	241
6.4 After the Bidding Ends	243
6.4.1 Bankruptcy and Non-performance	243
6.4.2 Scoring Rules vs. Price-Only Bids	245
6.5 Conclusion	247

PART II MULTI-UNIT AUCTIONS	251
7 Uniform Price Auctions	255
7.1 Uniform Price Sealed-Bid Auctions	257
7.1.1 Demand Reduction	258
7.1.2 Low-Price Equilibria	262
7.2 Simultaneous Ascending Auctions	265
7.2.1 The Simultaneous Ascending Auction and the Walrasian <i>Tatonnement</i>	268
7.2.2 Clock Auctions	279
7.2.3 Strategic Incentives in Uniform Price Auctions	284
7.2.3.1 The Basic Clock Auction Model	284
7.2.3.2 The Alternating-Move Clock Auction	287
7.2.3.3 Strategic Incentives with Elastic Supply	290
7.3 Conclusion	293
8 Package Auctions and Combinatorial Bidding	296
8.1 Vickrey Auctions and the Monotonicity Problems	302
8.1.1 Bidders' Vickrey Payoffs Bound Their Core Payoffs	305
8.1.2 Vickrey Auctions and the Entry Puzzle	305
8.1.3 When Are Vickrey Outcomes in the Core?	307
8.1.4 Substitute Goods and Core Outcomes	308
8.1.5 Substitute Goods and Vickrey Outcomes	312
8.2 Bernheim–Whinston First-Price Package Auctions	315
8.2.1 Formulation	316
8.2.2 Profit-Target Strategies	318
8.2.3 Equilibrium and the Core	319
8.3 Ausubel–Milgrom Ascending Proxy Auctions	324
8.3.1 The Proxy Auction with Unlimited Budgets	325
8.3.1.1 Proxy Outcomes Are Core Outcomes	326
8.3.1.2 Profit-Target Strategies and Equilibrium	327
8.3.1.3 The Proxy Auction When Goods Are Substitutes	329
8.3.2 The Non-transferable-Utility Proxy Auction	330
8.4 Conclusion	333
<i>Bibliography</i>	339
<i>Author Index</i>	347
<i>Subject Index</i>	351

Getting to Work

The era of putting auction theory to work began in 1993–1994, with the design and operation of the radio spectrum auctions in the United States. Although the economic theory of auctions had its beginnings in the 1960s, early research had little influence on practice. Since 1994, auction theorists have designed spectrum sales for countries on six continents, electric power auctions in the United States and Europe, CO₂ abatement auctions, timber auctions, and various asset auctions. By 1996, auction theory had become so influential that its founder, William Vickrey, was awarded a Nobel Prize in economic science. In 2000, the US National Science Foundation's fiftieth anniversary celebration featured the success of the US spectrum auctions to justify its support for fundamental research in subjects like game theory. By the end of 2001, just seven years after the first of the large modern auctions, the theorists' designs had powered worldwide sales totaling more than \$100 billion. The early US spectrum auctions had evolved into a world standard, with their major features expressed in all the new designs.

It would be hard to exaggerate how unlikely these developments seemed in 1993. Then, as now, the status of game theory within economics was a hotly debated topic. Auction theory, which generated its main predictions by treating auctions as games, had inherited the controversy. At the 1985 World Congress of the Econometric Society, a debate erupted between researchers studying bargaining, who were skeptical that game theory could explain much about bargaining or be useful for improving bargaining protocols, and those investigating in auctions and industrial organization, who believed that game theory was illuminating their studies. Although game theory gained increasing prominence throughout the 1980s and had begun to influence the

leading graduate textbooks by the early 1990s, there was no consensus about its relevance in 1994, when the Federal Communications Commission conducted the first of the new spectrum auctions.

The traditional foundations of game theory incorporate stark assumptions about the rationality of the players and the accuracy of their expectations, which are hard to reconcile with reality. Yet, based on both field data and laboratory data, the contributions of auction theory are hard to dispute. The qualitative predictions of auction theory have been strikingly successful in explaining patterns of bidding for oil and gas¹ and have fared well in other empirical studies as well. Economic laboratory tests of auction theory have uncovered many violations of the most detailed theories, but several key tendencies predicted by the theory find significant experimental support.² Taken as a whole, these findings indicate that although existing theories need refinement, they capture important features of actual bidding. For real-world auction designers, the lesson is that theory can be helpful, but it needs to be supplemented by experiments to test the applicability of key propositions and by practical judgments, seasoned by experience.

Whatever the doubts in the academy about the imperfections of game theory, the dramatic case histories of the new auctions seized public attention. An article in 1995 in the *New York Times* hailed one of the first US spectrum auctions³ as “The Greatest Auction Ever.”⁴ The British spectrum auction of 2000, which raised about \$34 billion, earned one of its academic designers⁵ a commendation from the Queen and the title “Commander of the British Empire.” In the same period, game theorists were plying their trade on another important application as well. The National Resident Matching Program, by which 20,000 US physicians are matched annually to hospital residency programs, implemented a new design in 1998 with the help of the economist–game theorist Alvin Roth. By the mid-nineties, thirty-five years of theoretical economic research about fine details of market design was suddenly bearing very practical fruits.

¹ See Hendricks, Porter, and Wilson (1994).

² See Kagel (1995).

³ The design was based on suggestions by Preston McAfee, Paul Milgrom, and Robert Wilson.

⁴ William Safire, “The Greatest Auction Ever,” *New York Times*, March 16, 1995, page A17, commenting on FCC auction #4.

⁵ The principal designers were Professors Ken Binmore and Paul Klemperer. They give their account of the auction in Binmore and Klemperer (2002). It was Binmore whom the Queen of England honored with a title.

1.1 Politics Sets the Stage

To most telecommunications industry commentators, the main significance of the US spectrum auctions was that a market mechanism was used at all. Spectrum rights (licenses) in the United States and many other countries had long been assigned in *comparative hearings*, in which regulators compared proposals to decide which applicant would put the spectrum to its best use. The process was hardly objective: it involved lawyers and lobbyists arguing that their plans and clients were most deserving of a valuable but free government license.⁶ With its formal procedures and appeals, a comparative hearing could take years to complete. By 1982, the need to allocate many licenses for cellular telephones in the US market had overwhelmed the regulatory apparatus, so Congress agreed to allow licenses to be assigned randomly among applicants by lottery.

The lottery sped up the license approval process, but it created a new set of problems. Lottery winners were free to resell their licenses, encouraging thousands of new applicants to apply for licenses and randomly rewarding many with prizes worth many millions of dollars. Lottery winners were often simple speculators with no experience in the telephone industry and no intention of operating a telephone business. Economic resources were wasted on a grand scale, both in processing hundreds of thousands of applications and in the consequent need for real wireless operators to negotiate and buy licenses from these speculators. The lotteries of small licenses contributed to the geographic fragmentation of the cellular industry, delaying the introduction of nationwide mobile telephone services in the United States.

A better process was needed, and in 1993, Congress authorized auctions as the answer. The question of how an auction market for radio spectrum should be designed was left to the Federal Communications Commission (FCC).

1.2 Designing for Multiple Goals

Congress did provide some instructions to the FCC governing the new spectrum auctions. One was that the first auctions were to be begun by July 1994. A second called for the auctions to promote wide participation in the new industry. The FCC initially responded to the second

⁶ The process was once characterized by an FCC Commissioner as “the FCC’s equivalent of the Medieval trial by ordeal” (as quoted by Kwerel and Felker (1985)).

mandate by introducing bidding credits and favorable financing terms for small businesses and woman- and minority-controlled businesses, to reduce the cost of any licenses acquired by those businesses. The statute also specified that the auction process should promote “efficient and intensive use” of the radio spectrum, in contrast with the fragmented use promoted by the lottery system. The meaning of the word “efficient” was initially subject to debate, but it was eventually read in economic terms to mean, in the words of Vice President Albert Gore, “putting licenses into the hands of those who value them the most.”⁷

There is a powerful tradition in economics claiming that individuals and firms, left to their own devices and operating in a sound legal framework, tend to implement efficient allocations. The argument is that when resources are allocated inefficiently, it is possible for the parties to get together to make everyone better off. So, following their mutual interests, the parties will tend to eliminate inefficiencies whenever they can. This traditional argument has its greatest force when the parties can all see what is required and have no trouble negotiating how to divide the gains created by the agreement. For radio spectrum, with thousands of licenses and hundreds of participants involved, computing just one efficient allocation can be an inhumanly hard problem, and getting participants to reveal the information about their values necessary to do that computation is probably impossible. Compared to the development of a universal standard (GSM) for mobile telephones in Europe, the more fragmented system that emerged in the United States highlights that the lottery system did not lead to efficient spectrum allocations. With so many parties and interests involved, the market took many years to recover from the initial fragmentation of spectrum ownership. During those years, investments were delayed and consumer services degraded. Getting the allocation right the first time does matter. Achieving that with an auction system called for a different and innovative approach.

The FCC, which the law had charged with designing and running the spectrum auctions, had no previous auction experience. Within the FCC, the design task was assigned to a group led by Dr. Evan Kwerel, an economist and long-time advocate of using auctions to allocate spectrum licenses.⁸

⁷ Quoted from Vice President Gore’s speech at the beginning of FCC auction #4.

⁸ Kwerel’s initial advocacy is explained in Kwerel and Felker (1985).

Like any other important FCC decision, the auction design decision would need to be based on an adequate public record – a requirement that would force the FCC to go through a long series of steps. It would need to write and issue a proposed rule, allow a period for Comments and another for Reply Comments, meet with interested parties to discuss and clarify the points of disagreement, resolve those disagreements, issue a ruling, consider appeals, and finally run the auction. Steps like these often stifle innovation, but that is not what happened on this occasion. With no political guidance about what kind of auction to use, no in-house experts lobbying to do things their way, and no telecom with an historically fixed position about how an auction should be run, Dr. Kwerel had unusual freedom to evaluate a wide range of alternatives.

Kwerel drafted a notice that proposed a complex auction rule. Industry participants, stunned by the novel proposal and with little experience or expertise of their own, sought the advice of academic consultants. These consultants generated a flood of suggestions, and the FCC hired its own academic expert, John McMillan, to help them evaluate the proposed designs. In the end, Kwerel favored a kind of simultaneous ascending auction, based in large part on a proposal by Robert Wilson and me and a similar proposal by Preston McAfee. The Milgrom–Wilson–McAfee rules called for a simultaneous multiple round ascending auction.⁹ This is an auction for multiple items in which bidding occurs in a series of rounds. In each round, bidders make sealed bids for as many spectrum licenses as they wish to buy. At the end of each round the *standing high bid* for each license is posted along with the minimum bids for the next round, which are computed by adding a pre-determined bid increment, such as 5% or 10%, to the standing high bids. These standing high bids remain in place until superseded or withdrawn.¹⁰ An *activity rule* limited a bidder's ability to increase its activity late in the auction, thus providing an incentive to bid actively early in the auction. For example, a bidder that

⁹ The principal difference was that the Milgrom–Wilson design proposed the now standard features that bidding on all licenses would remain open until the end of the auction, with progress ensured by Milgrom's activity rule. McAfee's design had no activity rule, and ensured the progress of the auction by closing bidding on each license separately after a period with no new bids on that license.

¹⁰ A bidder that withdraws its bid pays a penalty equal to the difference, if positive, between the eventual sale price for the license and the amount of its withdrawn bid. If the eventual price exceeds its bid, then no penalty is payable.

has been actively bidding for ten licenses may not, late in the auction, begin bidding for eleven licenses.

The theory of simultaneous ascending auctions is best developed for the case when the licenses being sold are substitutes. During the course of the auction, as prices rise, bidders who are outbid can switch their demands to bid for cheaper licenses, allowing effective arbitrage among substitute licenses. One of the clearest empirical characteristics of these auctions is that licenses that are close substitutes sell for prices that are also close – a property that is not shared by most older auction designs.

The initial reception to Kwerel's recommendation was skeptical. The proposed auction was unexpectedly complicated, and FCC Chairman Reed Hundt sought the advice of other FCC staff. He asked the economics staff: If you could pick any design you want, would this be it? He asked those who would have to run it: Can this really work? Even in the short time available to set it up? With the endorsement of his staff, Chairman Hundt decided to take the risk of adopting a new auction design.

1.2.1 Substitutes and Complements

Auctions are processes for allocating goods among bidders, so the challenge of auction design can only be understood by studying the demands of the participants. In the initial PCS auction, there were three groups of potential bidders. The first group included long-distance companies with no existing wireless businesses. These companies, including MCI and Sprint, were making plans to enter the wireless business on a national scale. Each wished to acquire a license or licenses that would cover the entire United States, allowing it to make its service ubiquitous and to combine wireless with its own long distance service to offer an attractive and profitable package to consumers.

A second group comprised the existing wireless companies, including AT&T, some regional Bell operating companies, and others. The companies in this group already owned or controlled licenses that enabled them to offer services to parts of the country. Their objectives in the auction were to acquire licenses that filled in the varying gaps in their existing coverage and to expand to new regions or perhaps the entire nation. These companies posed a regulatory challenge for the FCC, which wanted to allow them to meet their legitimate business needs without gaining control of enough of the spectrum to manipulate market prices. To avoid this outcome, the FCC imposed limits on the amount of

spectrum that any company could control in any geographic area. These existing wireless operators would be ineligible to bid for a nationwide PCS license of the sort that had typically been awarded in European countries. From MCI's perspective, this meant that a nationwide license might be bought cheaply at auction, so it lobbied the FCC to structure the new licenses in this way.

The last group consisted mainly of new entrants without wireless businesses. Some of these companies, like Pacific Bell in California, were quite large. These companies typically sought licenses or packages covering large regional markets, but not licenses covering the entire nation.

One of the first lessons to take from this description is that the auction game begins long before the auction itself. The scope and terms of spectrum licenses can be even more important than the auction rules for determining the allocation, because a license can directly serve the needs of some potential bidders while being useless to others. For the actual PCS auctions, a license provided its owner the right to transmit and receive radio signals suitable for mobile telephone service in a particular band of radio frequencies and in a particular geographic area. These license specifications constrained the possible spectrum allocations. For example, suppose three separate licenses covering areas A, B, and C were put up for sale. If one bidder wanted a license covering A and half of B while the other wanted a license covering C and the other half of B, the license specifications would prevent each bidder from acquiring its optimal allocation. One task of the auction designer was to promote the best (most "efficient") possible allocation, subject to such constraints.

Achieving efficiency involves various subtle complications. A certain license may be valuable to one bidder because it helps exclude entry and increase monopoly power, but be valuable to another because the buyer will use it to create valuable services. In comparing the efficiency of allocations, only the second kind of value counts, but bidders do not respect that difference when placing their bids. The value of a license to a bidder may depend not only on the license itself, but also on the identities of other licensees and the technologies they use. For example, the licensee identities can affect their "roaming arrangements" – which allow their customers to use another company's services when they roam to the other company's license area. A third complication is

that the bidders may need to pool information even to determine their own likely profits from various arrangements, for example because the bidders have different information about the available technology or forecasted demand.

But the fundamental barrier to efficiency that was most debated among the FCC auction designers concerned the *packaging problem*. The value of a license to a bidder is not fixed; it generally depends on the other licenses the bidder receives. For example, a bidder might be willing to pay much more per license for a package of, say, five or six licenses than for smaller or larger packages.¹¹ Until such a bidder knows all of the licenses it will have, it cannot say how much any particular license is worth.

Consider a situation with just two licenses. If acquiring one license makes a bidder willing to pay less for the second, then the licenses are *substitutes*. If acquiring one makes the bidder willing to pay more for the second, then the licenses are *complements*. With more than two licenses, there are other important possibilities, and these add considerable complexity to the real auction problem. For example, if there are three licenses – say A, B, and C – and a certain bidder anticipates needing exactly two of them to establish its business, then A and B are complements if the bidder has not acquired C, but they are substitutes if the bidder has already acquired C. Nevertheless, most economic discussions of the auction design are organized by emphasizing the two pure cases.

Recent auctions devised by economic theorists differ from their predecessors in the ways they deal with the problems of substitutes and complements. Our later analyses will show that some of the new designs deal effectively with cases in which the items to be traded are substitutes, but that all auctions perform significantly worse when licenses might either be substitutes or complements. The impaired performance may take various forms including a loss of efficiency of the outcomes, uncompetitively low revenues to the seller, vulnerability to collusion, complexity for the bidders, and long times to completion.

¹¹ An instance of this sort arose in the Netherlands spectrum auction in 1998, where most of the licenses were for small amounts of bandwidth. New entrants were expected to need five or six such licenses to achieve efficient scale and make entry worthwhile.

To illustrate how value interdependencies affect proper auction design, we turn to a case study in which the matter received too little attention.

1.2.2 New Zealand's Rights Auction

New Zealand conducted its first auctions of rights to use radio spectrum in 1990. Some of the rights took the traditional form of *license rights* to use the spectrum to provide a specific service, such as the right to broadcast television signals using those frequencies. Others consisted of *management rights* according to which the buyer may decide how to use the spectrum, choosing, for example, television broadcasts, wireless telephones, paging, or some other service. In theory, when management rights are sold, private interests have an incentive to allocate spectrum to its most profitable uses, but the problem of coordinating uses among licensees can also become more complex.

Acting on the advice of a consulting firm – NERA – the New Zealand government adopted a *second-price sealed-bid auction* for its first four auction sales. As originally described by Vickrey (1961), the rules of the second-price auction are these: Each bidder submits a sealed bid. Then, the license is awarded to the highest bidder for a price equal to the *second* highest bid, or the reservation price if only one qualifying bid is made. The auction gets its name from the fact that the second highest bid determines the price.

The idea of a second-price sealed-bid auction strikes many people as strange when first they hear about it, but on closer analysis, the auction is not strange at all. In fact, it implements a version of the ascending (English) auction¹² similar to the one used at Amazon.¹³

In an ascending auction, if a bidder has a firm opinion about what the item is worth, then he can plan in advance how high to bid – an amount that we may call the bidder's *reservation value*. At sites like eBay and

¹² The most common form of an ascending (English) auction is one in which the auctioneer cries out increasing bids and the bidders drop out when they are no longer willing to pay above the current price. The auction ends when there is just one remaining bidder. As the winning bidder is required to pay the current high price, it is optimal for each bidder to stay in the auction only until the current price is equal to his valuation ("reservation value") of the item and not thereafter.

¹³ eBay also runs a similar auction, but its fixed ending time involves additional gaming issues as described by Roth and Ockenfels (2000).

Amazon, the bidder can instruct a *proxy bidder* to carry out a *reservation value strategy*. The proxy keeps beating the current highest bid on the bidder's behalf so long as that bid is less than the specified reservation value. If everyone bids that way, then the outcome will be that competition ends when the price rises to the second highest reservation value, or thereabouts (with differences due to the minimum bid increment). If everyone adopts such a reservation value strategy, then the ascending auction is almost the same as a second-price auction.

Strategic considerations in a second-price auction are easy: each bidder should set his reservation value to what the object is worth to him. If it happens that the highest bid among the other bidders is greater than this value, then he cannot do better than to bid his reservation value, because there is no bid he could make that would win the auction profitably. If, instead, it happens that the highest competing bid is less than his value, then setting his reservation value in this way wins and fixes the price at what the competitor bid, which is the best outcome that any bid could achieve. Thus, regardless of the bids made by others, setting a reservation value equal to the bidder's actual value always earns at least as much as any other bid.

The second-price sealed-bid auction has two advantages over most other designs. First, it duplicates the outcome of an ascending bid auction with small bid increments, but without requiring the bidders to be assembled together or to hire agents to represent them in their absence. Second, it presents each bidder with a simple strategic bidding problem: each merely has to determine his reservation value and bid it. This also means that there is no need for any bidder to make estimates of the number of other bidders or their values, for those have no bearing on a rational bidder's optimal bid.

The second-price auction has a simple extension to sales of multiple identical items, and it, too, can be motivated by considering a particular ascending auction. For example, suppose there is such an auction rule with seven identical items for sale, to be awarded to the seven highest bidders in an ascending outcry auction. Again, bidders might sensibly adopt reservation value strategies, bidding just enough to be among the top seven bidders and dropping out when the required bid finally exceeds the bidder's value. An analysis much like the preceding one then leads to the conclusion that the items will be awarded to the seven bidders with the highest values for prices approximately equal to the eighth highest

value. To duplicate that with a sealed-bid auction, the rule must award items at a uniform price equal to the highest rejected bid. In such an auction, the right advice to bidders is simple: “Bid the highest price you are willing to pay.” A similar uniform-price rule has sometimes been used in the sale of U.S. Treasury bills.¹⁴

In New Zealand, the government was selling essentially identical licenses to deliver television signals. On the advice of its consultants, it did not adopt this highest-rejected-bid rule, but chose instead to conduct simultaneous second-price sealed tender auctions for each license. New Zealand’s second-price rules would work well in one case only: when the values of the items were independent – neither substitutes nor complements. In the actual New Zealand auction, it would have been difficult to give bidders good advice. Should a bidder bid for only one license? If so, which one? If everyone else plans to bid for just one license and picks randomly, perhaps there will be some license that attracts no bids. Bidding a small amount for every license might then be a good strategy. On the other hand, if many spread around small bids like that, then bidding a moderate amount for a single license would have a high chance of success. With licenses that are substitutes or complements, independent auctions inevitably involve guesswork by the bidders that interferes with an efficient allocation.

The actual outcome of the first New Zealand auction is shown in Table 1. Notice that one bidder, Sky Network TV, consistently bid and paid much more for its licenses than other bidders. The Totalisator Agency Board, which bid NZ\$401,000 for each of six licenses, acquired just one license at a price of NZ\$100,000, while BCL, which bid NZ\$255,000 for just one license, paid NZ\$200,000 for it. Without knowing the exact values of various numbers of licenses to the bidders, it is impossible to be certain that the resulting license assignment is inefficient, but the outcome certainly confirms that the bidders could not guess one another’s behavior. If Sky Network, BCL, or United Christian had been able to guess the pattern of prices, they would have changed the licenses on which they had bid. The bid data shows little connection between the demands expressed by the bidders, the numbers of licenses they acquired, and the prices they eventually paid, suggesting that the outcome was inefficient.

¹⁴ The Treasury rule sets a uniform price equal to the lowest accepted bid.

Table 1. Winning Bidders on Nationwide UHF Lots: 8 MHz License Rights

Lot	Winning Bidder	High Bid (NZ\$)	Second Bid (NZ\$)
1	Sky Network TV	2,371,000	401,000
2	Sky Network TV	2,273,000	401,000
3	Sky Network TV	2,273,000	401,000
4	BCL	255,124	200,000
5	Sky Network TV	1,121,000	401,000
6	Totalisator Agency Board	401,000	100,000
7	United Christian Broadcast	685,200	401,000

Source: Hazlett (1998).

A second problem was even more embarrassing to New Zealand's government officials.¹⁵ McMillan (1994) described it as follows: "In one extreme case, a firm that bid NZ\$100,000 paid the second-highest bid of NZ\$6. In another the high bid was NZ\$7 million and the second bid was NZ\$5,000." Total revenue, which consultants had projected to be NZ\$250 million, was actually just NZ\$36 million. The second-price rules allowed public observers to get a good estimate of the winning bidders' profits, some of which were many times higher than the price. To avoid further embarrassment, the government shifted from the second-price sealed-bid format to a more standard *first-price* sealed-bid format, in which the highest bidder pays the amount of its own bid. As we will see later in this book, that did not guarantee higher prices. It did, however, conceal the bidders' profits from a curious public.

The change in auction format still failed to address the most serious auction design problems. Unlinked auctions with several licenses for sale that may be substitutes or complements force a choice between the risks of acquiring too many licenses and of acquiring too few, leaving a guessing game for bidders and a big role for luck. Allocations are unnecessarily random, causing licenses to be too rarely assigned to the bidders who value them the most.

¹⁵ For a detailed account, see Mueller (1993).

1.2.3 Better Auction Designs

In the New Zealand case, alternative auction designs could have performed much better. For example, the government could have mimicked the design of the Dutch flower auctions. The winner at the first round would be allowed to take as many lots as it wished at the winning price. Once that was done, the right to choose next could be sold in the next auction round, and so on. No bidder would be forced to guess about which licenses to bid on with such an auction. Each bidder could be sure that, if it wins at all, it will win the number of lots or licenses anticipated by its business plan at the bid price it chose.

There are other designs, as well, that limit the guesswork that bidders face. A common one in US on-line auctions allows bidders to specify both a price and a desired quantity. The highest bidders (or, in case of ties, those who bid earliest) get their orders filled in full, with only the last winning bidder running the risk of having to settle for a partial order. As with the Dutch design, efficiency is enhanced because bidders do not have to ponder over which licenses to bid on, and such rules reduce the *exposure* risk that a bidder may wind up acquiring licenses at a loss, because it buys too few to build an efficiently scaled system.

1.2.4 The FCC Design and Its Progeny

In the circumstances of the FCC's big PCS auction, it was obvious that some licenses would be substitutes. For example, there would be two licenses available to provide PCS service to the San Francisco area. Because the two licenses had nearly identical technical characteristics and because, for antitrust reasons, no bidder would be allowed to acquire both, these licenses were necessarily substitutes. The argument that some licenses were complements was also made occasionally, but the force of the argument was reduced by the large geographic scope of the licenses.¹⁶

As in the New Zealand case, the main design issue was to minimize guesswork, allowing bidders to choose among substitute licenses based

¹⁶ Dr. Mark Bykowsky of the National Telecommunications and Information Administration (NTIA) was a forceful advocate of the view that licenses could be complements and proposed a complex package auction design to accommodate the possibility. His case that complementarity was important is more convincing for the later auctions in which smaller licenses were sold. Nonetheless, the short time available to run the first auction led to a near-consensus that the package auction proposal involved too many unspecified details and unresolved uncertainties for it to be evaluated and adopted immediately.

on their relative prices. When substitute goods are sold in sequence, either by sealed bids or in an ascending auction, a firm bidding for the first item must guess what price it will have to pay later if it waits to buy the second, third, or fourth item instead. Incorrect guesses can allow bidders with relatively low values to win the first items, leading to an inefficient allocation. With this problem in mind, the final rules provided that the licenses would be sold all at once, in a single open ascending auction, during which bidders could place bids on any of the licenses and track bids on all the licenses. The openness of the process would eliminate the guesswork, allowing bidders to switch among substitute licenses, and promote equal prices for licenses that are perfect substitutes.

In order for the auction to work in this idealized way, bidding on all licenses would need to remain open until no new bids were received for any license, but that raised a new issue. In a worst case scenario, the auction might drag on interminably as each bidder bid on just one license at a time, even when it was actually interested in eventually buying, say, 100 licenses. To mitigate this risk, the FCC adopted my *activity rule*. The general application of an activity rule involves two key concepts: eligibility and activity. A bidder's activity in any round is the *quantity* of licenses on which it has either placed new bids in the round or had the high bid at the beginning of the round. In the example cited earlier, the quantity is just the number of licenses on which a bid is placed, but other quantity measures, including the total bandwidth of the licenses bid or the bandwidth multiplied by the population covered, have also been used. The rule specifies that a bidder's total activity in a round can never exceed its eligibility. A bidder's initial eligibility, applicable to the first round of the auction, is established by filing an application and paying a deposit before the bidding begins. Its eligibility in each later round depends on its recent bidding activity. One simple form of the rule specifies that a bidder's eligibility in any round after the first is equal to its activity in the preceding round. Thus, bidders that are not active early in the auction lose eligibility to place bids later in the auction. This rule speeds the auction and helps bidders to make reliable inferences about the remaining demand at the current prices.

The FCC rules have evolved since the original 1994 design, but larger changes have been made to adapt the simultaneous ascending auction to other applications. One common variation arises when there are many