A Primer on Auction Design, Management, and Strategy*

David J. Salant TSE

31000 Toulouse

December 23, 2013 DRAFT

Contents

1	Intr	roduction	11
	1.1	Goals of this Primer	11
	1.2	What are Auctions?	14
		1.2.1 Why Auctions?	16
		1.2.2 Types of Auctions	17
	1.3	A New Age of Auctions	19
		1.3.1 New Types of Auctions	20
		1.3.2 Auctions Replacing Regulation	22
		1.3.3 Auctions in the Private Sector	24
	1.4	Why Auction Design (and Management) Matters	25
	1.5	Outline of This Primer	27
2	Gar	ne Theory, Auction Design, and Strategy	29
	2.1	Game Theory and Auctions	29
	2.2	Noncooperative Games	30
		2.2.1 One-Shot Auctions	30

		2.2.2	Normal Form Games	31
		2.2.3	Equilibrium in Normal Form Games	32
		2.2.4	Mixed Equilibrium in Spectrum Auctions: Some Examples	33
		2.2.5	Equilibrium in Multi-attribute Auctions	34
	2.3	Multi-	stage and Sequential Auctions	36
		2.3.1	Subgame Perfect Equilibrium	36
		2.3.2	Signaling Games	38
	2.4	Repea	ted Games	39
		2.4.1	Finitely Repeated Games	40
		2.4.2	Infinitely Repeated Games	41
		2.4.3	Overlapping Generations of Players	42
	2.5	Summ	ary	42
3	Rev	venue I	Equivalence	45
	3.1	The F	our Basic Auction Types for Single-Object Auctions	45
	3.2	Auctio	on Strategy in the Four Basic Auction Types	47
	3.3	Strate	gic Equivalence	48
	3.4	Reven	ue Equivalence of English and Dutch Auctions	51
	3.5	Summ	ary	55
4	Opt	imal a	nd VCG Auctions	59
	4.1	Optim	al Auctions	59
	4.2	Calcul	lating Optimal Bid Functions	60

	4.3	Optim	al Auctions	62
		4.3.1	The Revelation Principle	62
		4.3.2	The Revelation Principle and Optimal Auctions $\ . \ . \ .$	63
	4.4	Incent	ive-Compatible Auctions	66
		4.4.1	The Vickrey-Clark-Groves Mechanism	66
		4.4.2	Properties of VCG Auctions	68
		4.4.3	Multi-object VCG Auctions	69
		4.4.4	VCG-like Mechanisms for Advertising Auctions	72
	4.5	Summ	ary	73
5	Imp	erfect	Information and the Winner's Curse	77
	5.1	Comm	non Values and Order Statistics	77
	5.2	Biddir	ng Strategy in Common-Value Auctions	79
		5.2.1	Bid Strategy in Dutch and First-Price Auctions	80
		5.2.2	Bid Strategy in a Second-Price Auction	81
		5.2.3	Bid Strategy in English Auctions	81
		5.2.4	Comparisons between English, Second-Price, and Dutch Common-Value Auction Outcomes	83
	5.3	Almos tries	t Common-Value Auctions and Informational Asymme-	84
	5.3	Almos tries 5.3.1	Auctions with One Informed Bidder	84 85
	5.3	Almos tries 5.3.1 5.3.2	Auctions with One Informed Bidder Auctions with Both Informed and Uninformed Bidders	84 85 86
	5.3	Almos tries 5.3.1 5.3.2 5.3.3	at Common-Value Auctions and Informational Asymme- Auctions with One Informed Bidder Auctions with Both Informed and Uninformed Bidders Other Asymmetries	84 85 86 87

6	\mathbf{Seq}	uentia	l Auctions of Substitutes	91
	6.1	Multip	ple-Lot Auctions	91
	6.2	The D	Declining-Price Anomaly	93
	6.3	Weber	's Martingale Theorem	95
	6.4	Strate	gic Allocation Decisions	97
		6.4.1	A Simple Two-Auction Example	99
		6.4.2	Optimizing the distribution of auction target volume across auctions	103
		6.4.3	Sequential Auctions and Risk	106
		6.4.4	Auction Timing	107
	6.5	A Sun	amary of Auction Timing Considerations	108
7	\mathbf{Seq}	uentia	l Auctions of Complements	111
7	Seq 7.1	uential Exam	l Auctions of Complements ples of Complementarities	111 112
7	Seq 7.1 7.2	uential Examj Renew	l Auctions of Complements ples of Complementarities	111 112 114
7	Seq 7.1 7.2 7.3	uential Examp Renew Sequer	l Auctions of Complements ples of Complementarities	111112114116
7	Seq 7.1 7.2 7.3 7.4	uential Examp Renew Sequer Sequer	I Auctions of Complements ples of Complementarities val Auctions ntial Auctions of Complements ntial Intellectual Property Rights Auctions	 111 112 114 116 119
7	Seq 7.1 7.2 7.3 7.4 7.5	uential Exam Renew Sequer Sequer Auctio tion in	I Auctions of Complements ples of Complementarities val Auctions ntial Auctions of Complements ntial Intellectual Property Rights Auctions ons for Complementary Inputs with Imperfect Competi- n Downstream Markets	 111 112 114 116 119 122
7	Seq 7.1 7.2 7.3 7.4 7.5 7.6	uential Exam Renew Sequer Sequer Auctio tion in Summ	I Auctions of Complements ples of Complementarities val Auctions ntial Auctions of Complements ntial Intellectual Property Rights Auctions ons for Complementary Inputs with Imperfect Competi- n Downstream Markets	 111 112 114 116 119 122 124
8	 Seq 7.1 7.2 7.3 7.4 7.5 7.6 Sing 	uential Examp Renew Sequen Sequen Auctio tion in Summ gle-Pro	I Auctions of Complements ples of Complementarities val Auctions ntial Auctions of Complements ntial Intellectual Property Rights Auctions ons for Complementary Inputs with Imperfect Competi- n Downstream Markets nary oduct Auctions	 111 112 114 116 119 122 124 127
8	 Seq 7.1 7.2 7.3 7.4 7.5 7.6 Sing 8.1 	uential Examp Renew Sequer Sequer Auction tion in Summ gle-Pro Introd	I Auctions of Complements ples of Complementarities val Auctions ntial Auctions of Complements ntial Intellectual Property Rights Auctions ons for Complementary Inputs with Imperfect Competi- nownstream Markets oduct Auctions uction	 111 112 114 116 119 122 124 127 127

6

	8.3	Clock Auctions
		8.3.1 Clock Auctions and Sealed-Bid Auctions
		8.3.2 Clock Auctions, Sealed-Bid Auctions, and Market Power133
	8.4	Supply-Function Auctions
	8.5	Conclusions and Applications
9	Sim	ultaneous Auctions 139
	9.1	The SMR Auction
	9.2	Introduction to the SMR Auction
		9.2.1 A Brief History of the SMR and Clock Auctions 141
		9.2.2 SMR Auction and its Variants
	9.3	SMR Auction Properties
		9.3.1 Straightforward Bidding in SMR Auctions 144
		9.3.2 Strategic Withholding in SMR Auctions
	9.4	SMR Auction Experience
	9.5	SMR Auction Configuration
		9.5.1 The Texas Capacity Auctions
		9.5.2 The Spanish 4G Auction (2011)
		9.5.3 The Mexican PCS Auction
		9.5.4 US FCC Auction 66 – Advanced Wireless Services 158
		9.5.5 Arbitrage in SMR Auctions
	9.6	Conclusion

10 Con	binatorial auctions	165
10.1	Introduction	165
10.2	Hierarchical Package Bidding	168
10.3	VCG Auction Properties	170
	10.3.1 Revenues in VCG auctions	171
	10.3.2 Fairness in VCG auctions	172
	10.3.3 Budgets in VCG auctions	173
10.4	Core-Selecting Auctions	175
10.5	Ascending Package Auctions and the Combinatorial Clock Auction	177
	10.5.1 Complexity	179
	10.5.2 Bidder incentives	180
10.6	Comparing the Performance of the CCA and Other Auctions .	184
10.7	Conclusion	186
11 Fina	l Remarks	189
11.1	Introduction	189
11.2	Game theory and bidder incentives	189
11.3	Bidder incentives and revenue equivalence	190
11.4	Single-object vs multi-object auctions	191
11.5	Sequential auctions	192
11.6	Bidding behavior	192
11.7	Optimal auctions	193

8

CONTENTS

11.8 Auction design, management, and strategy	194
11.9 Participation \ldots	194
11.10Auctions and markets	195

CONTENTS

Chapter 1

Introduction

Summary This is a primer on auctions. It is intended to serve as a guide to auctions for the practitioner. This chapter explains what distinguishes auctions from other market transactions. It provides a brief summary of what is essential or helpful for the practitioner.

1.1 Goals of this Primer

Auctions are an increasingly common share of all market transactions. They take a great many forms. Often they greatly improve market liquidity, but often, in contrast, they totally fail to work. Perhaps more than half of all auctions result in some or all of what is being put for auction remaining unsold, and even when this does not occur, the transaction prices can be absurdly low or high, and far from equilibrium levels. This book explores reasons why auctions work or fail. It explains how auction design affects bidding decisions and how bidding decisions affect the outcome, sometimes quite dramatically.

Auctions are highly structured market transactions primarily used in thin markets (markets with few participants and infrequent transactions). In such markets, the standard supply-and-demand models of competitive market forces usually cannot be relied on to explain the outcomes, because in most auctions just a few bidders can have a large influence on the outcome. An auction is perhaps better characterized as a bargaining process. Moreover, in auctions, unlike most other markets, offers and counteroffers are typically made within a structure defined by a set of rigid and comprehensive rules. This book provides a fairly complete range of auction structures and discusses how they can affect auction outcomes.

This book is intended to serve both as an introductory text on auctions as well as a guide for practitioners, that is, those interested in managing or bidding in auctions. Auctions are highly structured negotiations with a defined set of rules for offers, counteroffers, price determination, and allocations. These rigid structures mean that auctions can be analyzed by mathematical models more accurately and completely than can most other types of market transactions.

This book provides a guide for modeling, analyzing, and predicting the outcomes of auctions. To be useful and effective, such a guide must cover some essential elements of auction theory. In practice, auction theory will rarely provide an exact and prescriptive model that can be applied directly. On the other hand, auction theory does provide both insights and specific results that are of direct value to the practitioner.

While a primer or practitioner's guide would ideally provide a complete checklist for what to do to design, set up, and manage an auction, or for bidders to decide how to bid, that scope is too broad to be practical. Rather, the focus here is on principles, tools, and examples that can be used to analyze auctions. This first requires some way to categorize different types of auctions. In this categorization, I also explain how the auction design and strategy should depend on the type of auction. Second, I explain the main results from game theory and the theory of auctions that can provide practical frameworks for analyzing auctions and bid strategy. Third, I review auction experience, mainly from actual auctions, but also from experiments. This part of the book illustrates how auction theory can be applied to auction design and strategy decisions.

Because of the increased use of auctions in new settings, auctions have become a more common subject of economic research and writing. Because this primer is intended as a guide for someone interested in designing an auction and managing it or for someone having to bid, its review of research on

1.1. GOALS OF THIS PRIMER

game theory and auction theory will stress the linkage between theory and application. Some game theory is crucial for the principles of sound auction design and for properly assessing bidding strategy.

This book largely avoids trying to provide formulaic and specific task plans or checklists, because most auction design and bid strategy problems are too complex to lend themselves to exact formulas. However, one of the main points of this book is to show how to apply game-theoretic modeling tools to practical auction design and strategy decisions. Practical difficulties are explained in some detail. Reality never exactly fits any theoretical model, but very simple theoretical models can still be very useful.

Explanations of how auction strategy and auction design can matter in one specific case may have only limited applicability to seemingly very similar specific cases. Similarly, the main results of auction theory may have only small direct practical value. What is often most useful are simple mathematical models and a few basic principles for developing them, as well as guidelines to help identify limitations of such models. *Simulations*, in which computer algorithms substitute for bidders, and scaled-down *experiments*, in which human subjects play the role of bidders in repeated trials, can also be useful tools for assessing auction design or bid strategy plans. Much of this book's exposition is devoted to explaining how to develop simple and useful models of auctions.

One additional objective in this book is to explain how both auction design and bid strategy can have a significant effect on prices and allocations. Auctions are a form of imperfectly competitive markets, in which both sides have limited information. More specifically, auction design can have a significant effect on revenues in a *forward* auction, that is, an auction to sell one or more lots, and on costs in a *reverse*, or procurement, auction. I provide many examples in which identical items have sold for significantly different prices in the same auction and at the same time. An analysis of how the auction design affects bidder incentives can usually go a long way in explaining such price anomalies. Price differences are even more common when identical objects are auctioned one at a time and in other types of sequential transactions.

I start with some basic exposition of game theory and with a few of the main results from auction theory. To be useful, a model of an auction must be consistent and must quantitatively represent the incentives and decisionmaking of the bidders. Mostly, this assumes rational bidders. However, behavioral and cognitive limitations of bidders are discussed.

This primer should not be considered as a comprehensive compendium of modern auction theory. Many topics are omitted - such as two-sided auctions or econometric analysis of auctions¹ - and there is relatively limited discussion of some more technical topics, such as mechanism design. The intent here is to provide an introduction and a starting point for the reader to either learn something about auctions or interested in auction design or strategy.² This primer still should also be of significant interest to researchers. There are many significant practical issues in auction design and strategy that have yet to attract significant attention from researchers. Some of the material in this book is new, largely motivated by a need to extend and expand the existing literature in order to address some significant auction design issues that have arisen in recent auctions.

1.2 What are Auctions?

This section briefly describes what auctions are, and the types of auctions studied in this book. I also provide a brief introduction to why auctions are used in preference to other market mechanisms.

To start, it is essential to say what constitutes an auction, or at least what will be included in this book. What distinguishes an auction from other types of transactions is not always clearly specified. I will use the term primarily to describe all one-to-many bidding processes: forward auctions, with one seller and many buyers, or reverse auctions, with one buyer and many sellers. At times, the term *auction* is extended to situations in which there are a few buyers or a few sellers, i.e., few-to-many transactions. But in these cases, the few will be assumed to be acting in concert in organizing the auction. When there are two sellers, acting totally independently, there are really two independent auctions. In contrast, when the two sellers agree on schedules, starting prices, or other elements of the bidding process, then the two sellers are really acting as a single entity in determining the rules of the auction.

 $^{^{1}}See [58].$

²See [46], [54] for more theoretically advanced treatments and also [41].

The term auction will be restricted here to what are often called *sealed-bid tenders*. More specifically, one-shot bidding processes, both price-only and multi-attribute ones, are referred to here as sealed-bid auctions. Most of the focus is on price-only or other forms of single-dimensional bidding.

More formally, an auction is defined by:

- 1. Bidding format rules, which govern the form of bids. Bids can be price only, multi-attribute, price and quantity, or quantity only.
- 2. Bidding process rules, including (a) closing rules, which determine the timing of the auction, (b) specification of the information provided to bidders during the auction, (c) rules about the processes for bid improvements or counter-bids, and (d) rules specifying the auction-closing conditions.
- 3. Price and allocation rules that determine final prices and the amounts won by each winning bidder. Auctions can be for single items or multiple items. Traditionally, multiple-lot auctions have been conducted as separate auctions, sometimes sequentially and sometimes simultaneously. Modern auction designs have been developed to facilitate bidder arbitrage for similar items or to allow bidders to submit offers for different packages of items.

An auction is one type of market mechanisms. More specifically, an auction is an allocation mechanism that mostly uses price criteria for determining *allocations* – that is, which parties win which objects or contracts – and the prices paid. Auctions are more centralized than most other market mechanisms. In most markets, other than *exchanges*, buy and/or sell offers are not all made in one place or submitted through an auctioneer. An auction, however, is not the only type of centralized market mechanism.

An auction is also a price discovery mechanism. Auctions provide a means for the entity conducting the auction – the auctioneer, or the auction *originator*, to collect information from bidders so as to determine prices and allocations. The fact that the information bidders provide can affect their prices means that it can also affect their incentives, as is explained in more detail below. An exchange (such as a stock, commodity, or financial exchange) is at times considered to be a form of auction. Exchanges differ from the auctions considered here in three respects. First, exchanges involve two-sided or manyto-many transactions. Second, transactions on an exchange normally run continuously, whereas auctions are periodic or episodic. Third, exchanges usually have an intermediary setting the price who is neither a net buyer nor a net seller on average. The intermediary (or specialist or trading desk) will have a book of buy and sell orders and adjust prices up or down depending on the relative volumes of such orders. In contrast, in auctions the originator will usually have an interest in getting the best price (i.e., the highest price in a forward auction, and the lowest price in a reverse auction). Auctions can, however, have independent entities to manage the bidding process.³

1.2.1 Why Auctions?

Auctions are not used for most transactions. The question arises why auctions are used at all. When the transaction frequency is high, transactions occur at physical or virtual stores, or in exchanges. When it is low, but the volume or monetary value is high, then an auction is most useful. When the transaction frequency is low, it will not be worthwhile to incur the cost of maintaining a trading site, or store. While the organization of an auction may involve some costs, when the trading events are infrequent, it can be cost effective for auctions to be used.

When the volume or the value is also low, then bilateral negotiations and the solicitation of offers through requests for offers (RFOs) or requests for proposals (RFPs) are often most efficient. When the transaction frequency is low, the cost of maintaining a store or exchange is high relative to the transaction value. When the transaction value is low, the cost of organizing and managing an auction can be prohibitive. Thus, auctions are best for periodic or episodic high-volume and high-value transactions.

³Sometimes, an auction originator will engage an independent third party to serve as an auction manager. However, the auction manager usually assumes an operational role. In contrast, the exchange specialist, or trading desk, will serve to adjust prices to balance supply and demand offers, and its operations are generally housed by an exchange that handles hundreds or thousands of types of assets.

1.2.2 Types of Auctions

Auctions have been around for centuries – at least since the time of Herodotus in 500 BCE. Auctions can take many different forms and can involve many types of goods and services. Two main ways in which auctions differ are in the auction format, or rules structure, and the frequency and volume of the transactions. Until recently, most auctions were single-object auctions or sequences of such auctions.

Auction Formats

Traditionally, most price-only auctions have used some variant on one of four basic forms:

- 1. English, ascending price
- 2. Dutch, descending price
- 3. first-price sealed bid
- 4. second-price sealed bid.

Some auctions allow bids to include multiple components or attributes, and not just price. In multi-attribute bidding processes explicit weights or subjective evaluations are used to decide winners.

The most common form of open auction is still open, oral-outcry *English* auctions. In a standard English auction, the auction manager or auctioneer announces prices – an increasing sequence of prices in a forward auction, and a decreasing sequence in procurement auctions – and bidders indicate whether they are willing to accept the announced price. Bidders need not indicate a willingness to accept each announced price in order to remain eligible, and once one bidder indicates acceptance of the most recent price, the auctioneer will go to the next increment or decrement. Bidders can often, if they desire, shout out (inject) jump bids. The auction ends when no one is willing to improve on the most recently announced price.

The most common form of *online* auction is still the English auction, or simple variants thereof. Instead of an auctioneer calling out prices, though, the host sets a starting price, and bidders can submit bids at any time during an open bidding window. Each new bid is posted, and the auction host will raise the minimum for subsequent bids, similarly to the manner in which a live auctioneer raises price in an English auction.

Descending-price forward auctions have been used since Babylonian times.⁴ The name "Dutch auction" to describe these types of auctions probably arose from their use in selling tulips in the Netherlands, which dates back at least to the 17th century.⁵ In a (forward) Dutch auction, the auctioneer starts with a high price, and gradually lowers the price until a bidder is willing to accept that price. Notice that for a reverse auction, that is, an auction in which bidders are sellers and the auctioneer is a buyer, a declining-price auction would elicit bids from all bidders at a high initial price. So, strategically, a declining-price forward (Dutch) auction, that is, an auction in which bidders are buyers, is strategically quite different from a declining-price forward auctions. The bid strategy is clearly different in a Dutch and an English auction. In a Dutch auction, a bidder has to guess when to bid. In an English auction, there is no such guessing. A bidder can stop when his or her value is reached.

Both first-price and second-price sealed-bid auctions, often called sealed-bid tender auctions, award to the high bidder the object for sale in a forward auction, and to the low bidder the contract to supply a product or service in a sealed-bid reverse auction. In a first-price auction, the winning bidder pays the bid amount in a forward auction, or is paid the bid amount in a reverse auction. In a second-price auction, the winning bidder pays, or is paid, the amount offered by the best losing bid.

Transaction Frequency and Volume

Auctions also differ in how many units are transacted and how often. An auction for a single object is a one-time event, and the strategic considera-

 $^{^{4}}$ See [76].

 $^{{}^{5}}See [26] and [70].$

tions are quite different than they are in multi-object auctions and repeated auctions. The theory and experience for one-shot auctions is a useful reference point for the subsequent analysis of multi-object auctions. Auctions can be one-time, or infrequent, events for a large number of objects. The major spectrum auctions in the US, Europe, and elsewhere are examples of such auctions. Such auctions have a *matching* role in determining which objects go to which bidders, as well as the usual allocation role in determining winners and losers. Bidders can also interact in the auction to try to more efficiently sort out who wins what, and at what prices. It is important that the auction rules allow efficient arbitrage, as otherwise both the auctioneer and the bidders can lose.

Auctions can also be recurring, such as the monthly capacity auctions conducted by many electricity transmission operators to ensure adequate generation availability. Many commodities – dairy products, wine, etc. – are auctioned periodically. Mostly the same bidders compete in each such auction, and for mostly the same types of products. These periodic auctions are commonly multiple-unit, and often multi-object. As such, they have many of the same strategic features as one-shot, multi-object auctions. However, the repetition allows bidders to communicate and interact over time across auctions. This can further affect bidding behavior and the outcomes of the auctions.

Auctions can be categorized in other ways. *Forward* auctions – that is, auctions in which there is one seller and several buyers – are perhaps the most common. But *reverse* auctions, in which there is one buyer and several sellers, are also fairly common. The modeling and analysis of forward and reverse auctions tend to be very similar.

1.3 A New Age of Auctions

Auctions have changed. They come in a great many new forms. A recent search at the US Patent Office website produced over 2,000 hits for patents issued that include the term "auction." While not every patent is for a different auction design, the number of innovations is still large. There are two main reasons for these changes, both in some sense owing something to John von Neumann (1903–1957), a mathematician who made pioneering contributions to both computer science and game theory.

One reason is the developments in game theory, a field that barely existed until the 1960s. Game theory is crucial, or should be crucial, to auction design. Game theory provides mathematical tools and techniques to analyze alternative auction designs. While game theory, and game theorists, are not always directly involved in auction design, their role has increased over the past ten to fifteen years, both in government agencies (such as the US Treasury, other national treasuries, and other governmental agencies in a number of sectors including telecommunications, energy, and natural resources) and in the private sector (where firms such as Google, Microsoft, Yahoo, HP, and IBM have hired leading game theorists to help improve their pricing tools: remember that an auction is a price discovery mechanism).

The other area in which von Neumann made pioneering contributions was in computer science. Without computers and the Internet, the use of auctions might still be where it was centuries ago. To say that auctions are a price discovery mechanism means bids provide the auction originator information about valuations, which are then used to determine prices and allocations. However, bidders will have incentives to withhold information if the information will affect prices paid. Iterative bidding, and other competitive processes, can mitigate bidders' incentives for strategic withholding. These iterative and competitive processes require efficient data exchange and processing facilities – facilities that would not otherwise be available without computers. Moreover, the use of the Internet greatly expands the pool of potential auction participants and reduces participation costs. These factors tend to improve auction results, and make auctions feasible for transactions that could never be realized absent computers and electronic communication networks.

1.3.1 New Types of Auctions

Over the past twenty years there has been a proliferation of different auction types. This section does not catalogue all of them. Instead, it briefly explains a few of the ways in which auction design has changed, due in part to changes in electronic communication technology, and in part to developments in game

1.3. A NEW AGE OF AUCTIONS

theory. The former allows more information to be transmitted and processed more quickly. The latter provides better incentives, and allows auctions to be introduced in areas where they may never have been possible before.

There have been significant advances in multiple-object auctions. A large class of simultaneous auctions for buying or selling a number of substitutes or complements at one time includes the *simultaneous multiple round* (SMR) *auction*, the *simultaneous ascending auction* (SAA), the *simultaneous descending clock auction* (SDCA),⁶ stage clock and package bid auctions, the Anglo-Dutch hybrid auction, simultaneous auctions with intra-round bidding, and many other designs. They have been introduced building off developments in game theory for auction design on the one hand, and in computer and communication technology on the other. In addition, a few types of package bidding have been introduced, including the SMR auction with hierarchical package bidding used for the \$19-billion auction of 700 MHz licenses conducted by the US Federal Communications Commission in 2008.

Even simple auction designs have changed. As noted above, online auctions typically employ a variant of the standard English auction in which a bidding window is open for a fixed duration, and bidders can enter offers as long as their offer exceeds the previous high bid by a minimum amount and the bidding window is open. This type of auction has been referred to as a "Yankee" auction. It has many variants, including one that extends the bidding window any time a new bid is entered. E-Bay has a "buy-it-now" feature that allows a bidder to enter a bid and close the auction.

Computers make possible package, or combinatorial, bidding. A package bid is an all-or-nothing bid on a combination of items. Package bidding is intended to address the *exposure* problem: a bidder having a value for a combination or package of items, but having a zero or sufficiently low value for the individual items, will be reluctant to submit bids on the individual items and risk winning them, unless that bidder is fairly certain it will win all the items in the package. Package bidding eliminates this exposure risk. However, package bidding adds a great deal of computational complexity because of the large number of possible combinations. When there are N

⁶The SMR and SAA auctions and the related SDCA allow for simultaneous auctioning in a sequence of rounds of multiple lots. Bidders topped in one round can respond in the next. Details are explained in Chapter 9 below.

items available, there will be 2^N possible subsets.⁷ This number will exceed 1,000 with only 10 items available.

Many new designs address the form of bids and the way bids are submitted to the originator.

1.3.2 Auctions Replacing Regulation

One area in which auctions have become an increasingly accepted transaction mechanism is in regulation, especially in the energy and telecommunications sectors. New Zealand and the US were the first countries to introduce auctions for allocating spectrum rights.⁸ Prior to auctions being used for allocating spectrum rights, a variety of allocation rules – including first-come, first-served; lotteries; and comparative hearings – were introduced. Auctions have now been introduced for allocating spectrum rights in dozens of countries. Spectrum auctions have perhaps been the largest auctions in history; the three largest spectrum auctions to date alone generated nearly \$100 billion in revenues – nearly \$50 billion for the German 3G auction, nearly \$35 billion for the UK 3G auction, and almost \$19 billion for the US 700 MHz auction.

Not all spectrum auctions have been successful. The success of the New Zealand auctions is debatable in view of their evidently low revenues and inefficiencies⁹. The US FCC's Auction #5 resulted in most of the largest bidders declaring bankruptcy shortly after the auction. These bidders tried to renegotiate lower prices, and a number of the large winners were successful in doing so. It took almost 10 years and a Supreme Court decision to resolve the conflict between bankruptcy law and the auction rules. Despite a few such problems, auctions appear to be a tool gaining increasing acceptance from regulatory agencies. To date the US Federal Communications Commission has conducted over 50 auctions, the UK, Australian, and Canadian telecommunication regulatory authorities have conducted over half a dozen auctions each, and telecommunication regulatory agencies in many other countries have used auctions for selling spectrum rights. Auctions therefore appear to

⁷This includes the empty set and the set itself.

 $^{^{8}}$ See [56] and [63].

⁹See Chapter 4 and [54] for a discussion.

1.3. A NEW AGE OF AUCTIONS

be a permanent part of the regulatory process for managing spectrum in a great many countries. Spectrum auctions will be discussed in more detail in Chapter 2.

Auctions are also increasingly common in the energy sector for electricity and gas. Simultaneous descending clock auctions are or have been used for energy procurement in a number of US states, starting with New Jersey and now including Montana, Illinois, Ohio, and California, and also in Spain and Italy. Other simultaneous auction formats are or have been used for selling entitlements for electricity in Belgium, France, Germany, Alberta (Canada), and Texas. Sealed-bid auctions are being used for energy procurement in Maryland, Delaware, Virginia, and the District of Columbia. Auctions are also being used for selling transmission or interconnector rights in the US and Europe, and for capacity transactions in much of the US. In the gas sector, various versions of the simultaneous multiple-round clock auction have been conducted in Austria, Germany, and Hungary for capacity rights.

Not all the experience has been positive. California has had particularly troubling experiences – in 1993, and then during the summer and Fall of 2000.¹⁰ The failure of California Power Exchange, a set of daily and longer-term auctions, is a well-documented example of how not to run an auction. The California QF auctions described below resulted in *negative* energy prices. And some Texas capacity entitlement auctions resulted in difficult-to-explain 50% or greater price differences for identical products in the same auction. Both of these auctions are described in more detail in Chapters 9 and 10.

In both the energy and telecommunication sectors, participation of qualified bidders is essential for a successful auction. Bidders don't always show up in sufficient numbers to make auctions competitive. Some examples in telecommunications include the first US auctions for 700 MHz spectrum and wireless communications services in 2000 and in 1997. In the latter case, the forecast revenues were about 100 times the actual auction revenues. In particular, the Congressional Budget Office had forecast a value of \$1.8 billion, but the auction raised only \$13.6 million.¹¹ Similarly, the Ohio–First Energy procurement auctions have failed over several years to attract adequate competition to allow First Energy to purchase any of its default service resources

 $^{^{10}}$ See [38] for a discussion, and [33].

¹¹See http://energycommerce.house.gov/comdem/legviews/mvbrspec.htm.

through the auction.

1.3.3 Auctions in the Private Sector

There may be more auction activity in the private sector than in the public sector, but much less information is publicly available about private sector auctions. Information about wholesale or business-to-business (B2B) transactions tends to be confidential. There has been an increasing reliance on auctions for both buying and selling in B2B dealings. Auctions are being used for everything from selling agricultural commodities to procuring professional services. A significant number of firms have arisen in the B2B sector, and many have folded or merged. Among the notable firms still around are Perfect Commerce (formerly Commerce One and Perfect), Ariba–Procuri (which includes what were formerly called Ariba, Procuri and Trading Dynamics), and Dovebid. Many of these online B2B auction firms have migrated into supply chain management, and many others have focused on one or two specific supply chains or "verticals," such as Hambricht for IPOs, ChemConnect for chemicals, and Nexant for energy.

For retail auctions, consumer-to-consumer (C2C) and business-to-consumer (B2C) E-Bay and Yahoo have the largest market shares. What is important for consumer auctions is participation, and so there has been some tendency toward increasing concentration on the consumer side. However, specialized sites such as those for travel and entertainment have been able to maintain some market presence. Consumer auctions tend to be less rigidly structured, be less promoted, and have lower participation than do B2B or government-run auctions. For this reason, less of what follows applies to these consumer auctions have encountered difficulties. There are a few exceptions, most notably the online auction site E-Bay and Google and Yahoo, which rely heavily on an auction pricing mechanism to price-position on their search pages.

1.4 Why Auction Design (and Management) Matters

What seems to be clear is that no one auction approach works all the time. Paul Klemperer states that auction design is a matter of the "horse for the courses" and not "one size fit all."¹² There are several reasons auction design matters.

First, not all auctions work equally well. There are many things that can go wrong. In some cases the possibility of misallocations caused by an auction design that puts bidders in a position where they are prone to make miscalculations and bad guesses tends not to matter as much when there are many competitors.

Second, in some cases, most notably in common value or affiliated value auctions, vigorous competition can result in overbidding and ex post performance problems.¹³

Third, in multiple-lot auctions, misallocations can occur for several reasons. A great deal of auction design work has been devoted to multi-product and multi-unit auctions.¹⁴ One of the main concerns in multiple-object auctions is the exposure problem, described above, where a bidder may have a low value for the individual items in a package on which it places a high value. Package bidding and simultaneous auction designs, including the SMR and the SMR with both hierarchical and nonhierarchical package bidding, were designed in part to address this concern. Auctions without package bidding are more likely to leave bidders unable to obtain efficient combinations of lots.

Package bidding raised other concerns, one of which is the *threshold problem*. When two (or more) bidders are seeking individual items for which a third bidder has placed a package bid, then the two bidders may need to find a way to coordinate their bids.

For example, suppose there are three bidders, B1, B2, and B3, and two items,

 $^{^{12}[41],\,\}mathrm{p.}$ 191.

 $^{^{13}\}mathrm{See}\ [17]$ for a discussion of oil lease bids with common values.

 $^{^{14}}$ See [64] for a discussion.

I and II. Suppose B1 wants item I, B2 wants II, and B3 wants both. If there is not much competition for I or II alone, neither B1 nor B2 might have to offer very much to be the high bidder on those items. Now, suppose that B3 has placed a bid of 8, which is posted, for both items together, and that B1 puts a value of 6 on item I and B2 puts a value of 6 on item II. If, absent B3, these two items would sell for 2 each, then B2 and B3 would each need to at least double their bids to win. Neither may want to do so without knowing how much the other might want to offer. This can result in B3 winning, when B1 and B2 should. The reverse can happen when B3 has the high combined value, but cannot enter a package bid. This possibility has been a topic of active research.¹⁵ and recently the UK government has introduced versions of *core-selecting* package auctions to address it. Chapters 9 and 10 presents some of the theory of SMR and package auctions. Chapters 9 and 10 also present experience with multiple-object auctions in the telecommunications and energy sectors. Package bidding also can add significant complexity. The number of possible combinations can get unmanageably large very fast.

One additional reason auction design matters is that some auctions are better at solving coordination problems for bidders.

On the flip side, there are instances in which it is possible for bidders to coordinate so as to divide the market and limit competition.¹⁶ As a response, regulatory authorities responsible for setting auction rules have limited the information reported to bidders during an auction. The better bidders are able to communicate, the easier is coordination, and at times such coordination can serve primarily to reduce competition in an auction. So the form of bids and the information reported to bidders can affect competition in an auction.

As much as design matters, how well an auction does also depends a great deal on implementation. It is important to emphasize that a primary requisite for any auction is well-prepared bidders. The rules, procedures, and management must facilitate bidder participation and not deter it. Too often the reverse is true, often because the originator wants to guarantee the best outcome from his or her perspective. By optimizing the bidding rules for the

 $^{^{15}}$ See [54].

 $^{^{16}}$ [18] document one case discussed below, and [15] explain the theoretical risks of open bidding. Also, see [50].

1.5. OUTLINE OF THIS PRIMER

originator, participation is discouraged, which paradoxically results in bad outcomes for the originator. And a competitive auction with less than ideal rules can often result in more efficient outcomes than will a good auction design with low participation. What is required is to enhance prospects for participation. What measures need to be taken so that participation does not sacrifice post-auction performance is discussed in Chapter 11.

However, auctions with multiple units of the same object do not always result in uniform prices. One of the great successes in auction design is the development of efficient auction designs for multi-unit auctions, which are easy to run. Clock auctions are quite efficient, and fast, for multi-unit, single-product auctions. Clock auctions and SMR auctions also work quite well when the products are substitutes and there are no large bidders. As has been noted elsewhere,¹⁷ however, clock and SMR auctions can be vulnerable to collusive or coordinated bidding. Klemperer designed the Anglo-Dutch hybrid to try to strike a balance between the benefits of an open SMR auction and the risks of collusion.

Many new auction designs involve improvements to the bid submission process or to accommodate online bidding. The time restrictions imposed in many online auctions, specifying a narrow bidding window, create incentives for bidders to wait until the last possible instant to bid – an activity called *sniping*. The clock auction eliminates this incentive, as the auction manager, rather than the bidder, raises the price. However, for reasons explained below, clock auctions are less practical for many consumer and B2B transactions. This has spawned a great deal of effort to adjust for sniping.

1.5 Outline of This Primer

This remainder of the primer is divided into three parts: (1) theory, (2) auction practice – that is, design, organization, and management, and (3) experience.

The theory part starts with a very brief introduction to essential concepts from game theory. The game theory framework is an essential tool for an-

 $^{^{17}}$ See [41] or [18] for a discussion.

alyzing auctions. I then explain a number of the more significant results from the auction theory literature, including the revenue, or payoff, equivalence theorem, the winner's curse, optimal auction design, and the theory of simultaneous and sequential auctions. The revenue equivalence theorem is worthy of specific mention as a key result that has widespread applicability. The analysis of simultaneous and sequential auctions is of great practical concern, as an auction originator with multiple objects, or multiple units, to auction will need to decide on whether to run one auction or more, and if more, how many.

The auction practice chapters address measures that can be taken to mitigate bidder collusion and to enhance competition, provisions for information disclosure, volume adjustments, and other basic principles of auction management that can enhance prospects for participation and a competitive outcome. Information disclosure provisions need to balance the benefits of information pooling and the risks of collusion. While full disclosure of all bids can invite market division and coordinated bidding, limited disclosure can often achieve the benefits of full disclosure without increasing the risk of anticompetitive bidding behavior. Volume adjustments are another new administrative tool that can be used to encourage more competitive bid-These issues are discussed in more detail below. This part of the ding. book also contains some discussion of auction experiments and simulations. Experiments are an increasingly popular research area in which subjects are recruited, and compensated, for participating in auctions staged in controlled environments. Simulations, on the other hand, rely on computer bidders and allow the computer to calculate the outcomes. Both forms of tests have benefits as well as limitations.

The third part of this primer discusses experience. Most of the discussion surrounds energy and telecommunications auctions, as there is enormous experience in these sectors, and they have been the proving ground for many new auction designs. These sectors also include some of the largest and most newsworthy auctions. This part will discuss auctions in other sectors, including natural resources, commodities, and general procurement.

Chapter 2

Game Theory, Auction Design, and Strategy

Summary

This chapter provides a brief introduction to game theory. Game theory provides analytical techniques needed to analyze auction design and strategy. This is a nontechnical introduction. The chapter focuses on how to apply game-theoretic techniques to the analysis of bid strategy and auction designs.

2.1 Game Theory and Auctions

This chapter provides a brief introduction to game theory. An appendix provides additional background on some of the more technical topics, especially those related to signaling games. Auctions have specific rules; these rules govern participation requirements, the form of bids, winner determinations, and payments. The rules of an auction can almost always be translated into a precise mathematical formulation. Such formulations are what is more formally called a mathematical game. A mathematical game is a representation of set of strategic interactions, similar to those that occur in ordinary games, into mathematical terminology. Such representations can facilitate analysis of the strategies and outcomes, or *equilibria*. This does not mean that it is always easy to *solve* an auction, as one would solve a system of equations. It not always easy, or even possible, to find or write a computer program even to compute auction outcomes. However, mathematical descriptions can lead to algorithms that solve part or all of an auction. Often it will be possible to capture the algorithms in computer programs or solutions of systems of equations, or to identify what specifically needs to be solved. Numerical techniques can also be applied to check for consistency of any planned strategy or desired goal of an auction design, using the underlying mathematical structure of an auction.

This chapter begins by describing what constitutes a *game*. Then some basic solution concepts for games are presented. It is explained how an auction can be interpreted as a game, and how the solution concepts from noncooperative game theory can be used to identify equilibria, likely outcomes, of auctions.

2.2 Noncooperative Games

An auction is generally best modeled as a *noncooperative* game. This does not mean that bidders will not cooperate in an auction; it only means that the bidders can have conflicting and/or competing interests. Auctions can be single-stage or multi-stage or multi-round. Single-stage and multi-stage games are represented in two distinct ways.

2.2.1 One-Shot Auctions

Many auctions are one-shot games. For example, in a sealed-bid auction, bidders name prices, and no bidder has any chance to respond. Sometimes, sealed-bid auctions also allow bids to include not only an offer but also desired quantities and nonprice attributes. Usually, the high bidder wins, and the amount paid is the bid amount. But other payment rules are common.

2.2. NONCOOPERATIVE GAMES

Row/column player	Left	Right
Тор	(1,1)	(0,3)
Bottom	(3,0)	(2,2)

Table 2.1: Prisoner's dilemma

2.2.2 Normal Form Games

A single-stage game is most often described in what is called *normal*, or strategic, *form*. A normal form game is defined by a triple: a set of players N = (1, 2, ..., N), a strategy set of each player, S_j , $j \in N$, and a payoff function for each player, $\pi_j(s_j, s_{-j})$, describing player or bidder j's payoff or profits as a function if its strategy s_j and that of all the other players, s_{-j} :

The prisoner's dilemma is a classic example of a noncooperative game. In it the row player's strategy set is {Top,Bottom} and the column player's strategy set is {Left,Right}. The payoffs are given by the entries in the above payoff matrix.

There are many types of these matrix games. In the prisoner's dilemma, the players have dominant strategies¹ of always playing the strategy that leads to a joint payoff minimum. Another basic example of a normal form game is a competitive game, e.g., the matching pennies game. In the matching pennies game each of two players can choose either heads or tails. If the choices match, then the first player wins both pennies, and otherwise the second player does. In this game, if one player does better, the other does worse, and that is why it is called a competitive game. In this game, there is no equilibrium in *pure* strategies. If the first player were to choose heads, then the second player would want to choose tails, and vice versa. Similarly, whatever is the choice of the second player, the first player would want to match that choice.

In order to characterize what players would do in this type of setting, one must consider random, or *mixed*, strategies. A mixed strategy for the row player in the prisoner's dilemma is a probability of choosing Top, and similarly a mixed strategy for the column player in the prisoner's dilemma is a probability of choosing Left. In the matching pennies game, if player 1

¹This concept of dominant strategies is discussed below.

were to choose heads with probability $\frac{1}{2}$, then player 2 would be indifferent between choosing heads and choosing tails, and similarly, if player 2 were to choose heads with probability $\frac{1}{2}$, then player 1 would be indifferent as to whether it chose heads or tails. This type of randomization, each player choosing heads with probability $\frac{1}{2}$, is then an equilibrium, as neither player would have an incentive to choose any other strategy.

A sealed-bid auction is a normal form game. In a first-price sealed-bid auction for a single lot, bidders' strategies are prices for the lot; the high bidder wins, receives the object, and pays the bid amount. The payoff for the high bidder is the value less the amount paid. All other bidders receive a zero payoff. A second-price sealed-bid auction is the same as a first-price auction, except that the high bidder only pays the second highest bid amount. As in the matching pennies game, some auctions will not have any pure strategy equilibrium but will have equilibrium in mixed strategies.

2.2.3 Equilibrium in Normal Form Games

Outcomes of auctions can often be readily predicted through equilibrium analysis of simple normal form models. The most basic form of equilibrium is called a Nash, or noncooperative, equibrium. A *Nash equilibrium* is a set of strategies, one for each player, such that no player can make a unilateral change in its strategy and increase its payoff. A Nash equilibrium is a minimum requirement for each entity to be acting consistently in line with each entity's view of its own interests.

One alternative, and equivalent, definition of a Nash equilibrium, for noncooperative games in general and a prisoner's dilemma in particular, is that each player has an expectation about what its rivals are doing, its choice of strategy is a best reply to those expectations, and every player's expectations are the same as the actual strategies chosen.

In the prisoner's dilemma example, the only Nash equilibrium is (Top, Left). This is clearly suboptimal for the players. Indeed, this outcome is a *dominant strategies equilibrium*: each player's choice of strategy is the same, no matter what the other chooses. A multi-unit auction can have a prisoner's dilemma component. Suppose, for example, that there are two blocks available, and

Col\row player	Heads	Tails
Heads	(+1,-1)	(-1,+1)
Tails	(-1,+1)	(+1,-1)

Table 2.2: Matching pennies

each of the two high-value bidders has declining marginal values, say H for the first block won and M < H for the second block won. Also suppose that other bidders have values of no more than L < M. Then, it will be an equilibrium for each firm to keep bidding in an ascending price auction up to M for a second block, even when their both stopping at L would result in higher surplus and the same allocation for the two high bidders.

In the matching pennies game, there is also a unique equilibrium: for each player to choose heads with probability $\frac{1}{2}$. This is a mixed strategy equilibrium. It is also not a dominant strategies equilibrium, as the optimal choice for one player depends on what other players choose:

2.2.4 Mixed Equilibrium in Spectrum Auctions: Some Examples

This section provides a few sample characterizations of Nash equilibria in some actual spectrum auctions. These cases are examples of actual auctions which only have mixed strategy equilibrium.

The first example is an auction of 2G spectrum that the Dutch government was planning to run in 1998. The government changed the rules after hearing bidder concerns. In that auction, there were 16 very similar blocks of spectrum for sale, and bidders were to be asked to submit a final round of sealed bids for each of those blocks. Bidders would need to win between 4 and 6 blocks in most cases to derive any value. A simplified version of this game was analyzed extensively by [67], in which two bidders are competing for three blocks and each has to win at least two to derive any value.² They

²To see that there may be no pure strategy equilibrium, suppose that each bidder has the same value v, and that each bidder j offers $a_j, b_j, (v - a_j, b_j)$ on blocks j = 1, 2, 3. If $a_j \ge b_j \ge 1 - a_j - b_j$, then the other bidder can bid a bit more than j on the second and third blocks and win both.

were able to completely characterize the equilibrium probability distribution for the mixed strategy of this auction.

Another example of auctions which will have no pure strategy equilibria comprises the *supply function* auctions commonly used in energy procurement. These auctions require bidders to submit a separate price for each block (a step function bid).³ The bidders' offers are aggregated to form a supply schedule, and the lowest offers that meet demand are the ones that are selected. In a *paid as bid* auction, bidders will receive their bid amounts for the portion of their supply that is selected by the auctioneer.⁴ The prices each bidder receives can be uniform, or paid as bid. In a uniform price auction all winning bidders receive the market-clearing price.

2.2.5 Equilibrium in Multi-attribute Auctions

At times, the auction originator's objectives will include nonprice factors. This is especially true when the auction is being run by a government agency which has a mandate to promote the public interest. Many spectrum auctions include coverage or buildout provisions. And at times, bidders receive a score based on price and nonprice submissions. For example, a recent French auction for 4G spectrum licenses weighted each bidder's monetary offer by factors that depended on whether the bidder would cover rural areas and provide wholesale access to resellers.⁵ The weights for the nonprice attributes,

 $^{^{3}}$ See [48] and [36].

⁴What follows is a very simple example. Suppose there are two suppliers in a region, each having capacity of 800 MW. This means that both firms are needed, and either firm can always supply 400 MW at the ceiling price of 100. Normalize price so that demand net of operating costs (assumed to be the same for both) is 1200 MW for $p \leq 100$ and zero otherwise, where p denotes the price. In a symmetric equilibrium, each firm will choose a probability distribution function F(p) which makes the rival indifferent as to what price it chooses. Neither firm will ever offer a price above 100 - or below 50, as for prices below 50 a firm can earn more by offering the ceiling price. Thus, 400pF(p) + 800p[1 - F(p)] = 4000, or $F(p) = 2p - \frac{100}{P}$, for $50 \leq p \leq 100$. For p < 50, we have F(p) = 0, and for p > 100, F(p) = 1. Each firm choosing its price with a probability determined by F(p) is an equilibrium in this example.

 $^{^5 {\}rm See} http://www.arcep.fr/fileadmin/reprise/dossiers/4G/proj-dec-appel-800mhz-160511.pdf and http://www.arcep.fr/fileadmin/reprise/dossiers/4G/proj-dec-appel-2-6ghz-160511.pdf.$

which were 0 or 1 decisions for the bidders, were high. This resulted in all winning bidders including these attributes in their offers.

Because bidders will often try to maximize net returns or profits, when both price and nonprice attributes contribute to the winning score, an equilibrium bid strategy will often be a *corner* solution; in other words, bidders will submit the maximum or minimum allowed values for the nonprice component of the bid, according as the bidder weighs the attribute more or less than the auctioneer.

One extreme example of this property of multi-attribute auctions was a California procurement of "qualifying facilities" (QF) in a sealed-bid auction. The auction was designed to ensure supply of renewable energy sources. Numerous factors were included in determining the score of each bidder.⁶ The auctioneer was seeking to procure electricity from renewable resources for both peak and off-peak periods, and to promote the construction of renewable generation capacity. Bidders were required to submit separate prices for different components of the generation service: one was a price required for ensuring capacity availability, and others included prices for energy supplied in r different periods. Wind farms were among the most significant sources of renewable energy. However, wind turbines only operate when there is wind. Therefore, the maximum capacity of a wind farm (or of other renewable energy source such as a hydroelectric facility) may not represent average availability. Thus, the auctioneer wanted to provide different weights on potential availability and actual production. The weight given to capacity in the score was relatively high; the result was that the winning bidders ended up receiving high capacity prices and *negative* energy prices. In other words, these wind farm developers received less money the more energy their wind farms supplied.

If the auctioneer can measure the monetary value of the nonprice attributes, then using those weights will provide bidders the appropriate incentives to submit optimal combinations of price and nonprice attributes. Simple Nash equilibrium analysis of bidder incentives when bids include price and nonprice attributes will usually suffice to anticipate the type of pitfalls that were experienced in the California QF auction.

⁶The details of this process are described in [33].

2.3 Multi-stage and Sequential Auctions

In many instances an auction will have multiple stages, so that bidders can respond to rivals' bids. In other cases, there will be a sequence of auctions for substitutes, for complements, or over time. Equilibrium criteria for one-shot games cannot allow for interactions and information revealed between stages or auctions. The *extensive form* representation of a game can be used to capture these effects.

An extensive form representation of a game includes three components: a game tree, decision nodes, and information sets. The tree describes the possible sequences of decisions and events, which can include random events. The decision nodes are points at which a player can take an action. And the information sets represent the information available to a player at each of the decision nodes. The graph of the prisoner's dilemma in Figure 2.3 shows how a normal form game can be represented in extensive form.⁷

Some auctions, such as the Federal Communications Commission (FCC) spectrum auctions, allow bidders to gradually improve their offers over the course of the auction. These multi-stage (or multi-round) auctions can be represented in extensive form.

Other times, for example in energy markets, the same bidders compete repeatedly over time. These type of auctions can be characterized as dynamic games, and can be represented in a type of extensive form, as is illustrated above.

2.3.1 Subgame Perfect Equilibrium

Multi-stage games differ from one-shot games in that the players interact repeatedly in the former and only once in the latter. After each round of interaction, the players can revise their plans. The notion of a subgame perfect equilibrium (SPE) is intended to ensure that equilibrium strategies remain so through the entire course of play. More specifically, a *subgame*

⁷The circle around the Column Player nodes in Figure 2.3 indicate that the column player does not know whether the row player has chosen "Top" or "Bottom" when it is choosing between "Left" and "Right".


Figure 2.1: Prisoner's Dilemma in Extensive Form

perfect equilibrium in an extensive form game is a Nash equilibrium with the additional property that the equilibrium strategies remain in equilibrium when restricted to any proper subgame, that is, any intermediate node of the game in which players know the prior history. What this means is that after each stage, round, or period, when a player has to make a decision knowing previous decisions of other bidders, then the bidder will not want to revise its strategy at that point.

As an example, consider an auction for three blocks, 1, 2, and 3; suppose further that there are two strong bidders, A and B. First suppose that the bids are collected all at once, but are opened one at a time. Also, suppose that a bid can be contingent on whether or not a bidder won one of the previous blocks.

Suppose too that the price for each block is the best losing bid for that block. Suppose both A and B prefer block 1, and that A is the only bidder who has a high value for block 3. Also, suppose that blocks 1 and 2 are very close substitutes, but bidder A is the stronger bidder. In this case, the following strategies are a Nash equilibrium, but not subgame perfect: Bidder A bids low for block 1, high for block 2, and high for block 3, contingent on winning one of blocks 1 and 2. Bidder B bids high for block 1, low for block 2, high for block 3 if it wins block 2, and low for block 3 if it wins only block 1.

B's high bid for block 3 in the event that it loses block 1 is a noncredible

threat to deter A from competing for block 1. What is a SPE is for B to bid low for block 1 and high for block 2, and vice versa for A. If B were to lose block 1, it would have no incentive to carry out its implicit threat.

Further stability requirements can, at times, further refine the notion of a SPE. 8

2.3.2 Signaling Games

In many multi-stage auctions, bidders can signal interests or intent.⁹ For instance, the United States FCC has conducted several dozen multi-round spectrum auctions. In these auctions, early round bids are often well below final prices, and do not represent any commitment by bidders, as the bids are low and certain to be topped in later rounds. In such auctions, concerns have been raised that bidders can signal so as to achieve coordinated outcomes.¹⁰ The question addressed here is to what extent one can model such behavior as an equilibrium outcome. Contrary to what may be popular opinion and some literature, collusive signaling need have no effect on the outcome. This is not to suggest that signaling in auctions has not occurred and/or has been ineffective. Rather, signaling is as likely to be ignored, or even counterproductive, as it is to be effective.

To see this, consider a simple situation in which there are two bidders, A and B, and two blocks, 1 and 2. Suppose too that each wants one block, and both prefer block 1, but both prefer block 2 to having to compete for 1. So, there can be two equilibrium outcomes: 1 goes to A and 2 to B, and the reverse. Now, suppose that before bidding each bidder can announce which block it will bid for. Presumably, if A announces 1 and B announces 2, then that should be the outcome. However, neither bidder may trust the other, and either can misinterpret the other. Absent an ability to form an

⁸A further requirement that can be imposed on equilibrium in a multi-stage game is that the equilibrium be *stable* for various forms of perturbations. Small variations in payoffs can at times disrupt some SPEs. Further small changes in strategies can also do the same. The ideas of a *trembling hand* perfect equilibrium (see Selten (**X**)) and the *stable set* of equilibria (see [45] provide further refinements of the SPE criterion.

 $^{{}^{9}}See [15] and [18].$

¹⁰See [15], [18], and www.fcc.gov/wtb/auctions.

ex ante binding agreement, these announcements are cheap talk, and can be ineffective or worse. 11

Indeed, one case that [18] cite about how bidders can signal each other in spectrum auctions shows how such efforts can backfire. In that auction, bidders sent signals – in the form of jump bids and overbidding – that two bidders were able to convey, observe, and (apparently) understand during the auction. However, one of the bidders, the *receiving party*, did not really want to coordinate in the way that the *sending party* signaled. During the auction, the sending party could only see that the receiving party was apparently acting cooperatively, and not that the receiving party had contacted the regulator. After the auction, the sending party was subject to prosecution and fined for attempting to collude during the auction.

2.4 Repeated Games

At times, the same bidders compete repeatedly for more or less the same lots. An example of this is the daily auctions conducted on most power exchanges.¹² The simplest form of repeated interaction occurs when the same game is repeated over and over. The notion of a SPE can be extended to these type of games. As infinitely repeated auctions have no last auction, rivals' bids in later auctions can depend on the outcome in a given auction, can pull a bidder in virtually any direction. This means, when the game is repeated indefinitely, there are infinitely many possible SPE outcomes.

In contrast, for many repeated games, including all those with a unique equilibrium in the one-shot stage game, there will be only one equilibrium in the finitely repeated game, no matter how long the interaction lasts, as long as there is a fixed end date and players are rational.

What follows is a brief description of the types of outcomes that can arise in repeated interactions. The discussion assumes that exactly the same game

¹¹Signaling games, such as those in which bidders send signals in auctions, typically have a multiplicity, and often a continuum, of equilibria. See [69] for a discussion of signaling games.

 $^{^{12}}$ See, for example, [38] for a description of one such market.

is repeated. However, most of the analysis applies to more general dynamic games.

2.4.1 Finitely Repeated Games

Consider the case in which the above prisoner's dilemma is repeated a fixed and finite number of times. Suppose that at each stage, or iteration, the row player chooses Top or Bottom, and the column player chooses Left or Right. These choices are made simultaneously. After the choices are made, the players are each informed about the other's choice, payoffs for the stage are accrued, and then the process repeats.

One could consider the possibility that each player will, in early stages, play cooperatively, and do so at least until near the end, as long as the other player does so. This is not a SPE, as each player can infer that the other will play noncooperatively (i.e., Top or Left) in the last period. Therefore, each player has an incentive to play Top or Left in the next to last period too.

The same logic will then apply to the third to last stage, as each player knows the other will choose to play noncooperatively in the second to last stage. By backward induction, there will be only one SPE in the entire game, which is for the row player to always choose Top and the column player Left.

If the game has a probability p of running an additional period, so that the end is uncertain, then other outcomes are possible. More specifically, if either player defects from a cooperative strategy, it will earn 3 for one period, and 1 for all subsequent periods – for however many periods the game lasts. The present value of this strategy is $3 + \frac{\delta p}{1-\delta p}$, where δ is the discount factor. In contrast, continued cooperation results in an expected payoff of $2 + \frac{2\delta p}{1-\delta p}$. And this quantity is larger than the first whenever $2\delta p > 1$, that is, when discounting is not too fast and the probability of the game ending fast not too high.

So, for example, in repeated interaction of traders in energy markets, there is a possibility of the traders all refraining from aggressive bidding, even where there is uncertainty. When there are multiple Nash equilibria in the stage game, cooperative outcomes can still be stable with finite interaction, provided that there are bad and good equilibria, and that each player wants to avoid the bad equilibrium in the last stage.¹³

2.4.2 Infinitely Repeated Games

An infinitely repeated game is one in which the same players interact in the same strategic situation indefinitely. The main, and big, difference between finitely repeated games and infinitely repeated games is that there is no last period. This means that, unlike a finitely repeated game, the backward unraveling from cooperative Nash equilibria does not occur. Therefore, cooperative outcomes are possible. Indeed, in infinitely repeated games, it can be shown that any payoff combination in the stage game that dominates the lowest payoff that each player can guarantee itself (its minmax payoff) can be supported using *grim trigger strategies* as a SPE, as long as the discounting is not too fast. A grim trigger strategy is one in which each player adheres to the cooperative play as long as no other player has deviated in any previous round. If there was a deviation by a player, then everyone reverts to that strategy that gives the defector its minmax payoff thereafter. When discounting is slow enough, the one-shot benefit from a deviation is outweighed by the benefits of continued cooperative play.¹⁴

This means that the outcome of repeated auctions can be indeterminate. As is discussed in more detail below, information that limits the ability of bidders to coordinate actions can induce bidding to be more competitive. This indeterminacy result, called a *folk theorem* for repeated games, can be extended to situations in which players cannot directly observe rivals' actions. When actions are not directly observable, payoffs are still observable, and so a player might become increasingly certain over time that a rival or rivals are deviating from cooperative play. Nevertheless, cooperative play can be a SPE when each player uses a stochastic decision rule: continue to cooperate as long as the price, or payoff, has averaged at least a threshold level, and if not, then revert to the Nash equilibrium (*punishment*) for a fixed and finite number of periods.¹⁵

¹³See [25] for a discussion of cooperative equilibria in finitely repeated games. ¹⁴See [2].

¹⁵For a more detailed discussion of cooperative equilibria with demand or cost uncertainty, see [30].

The reason to use a stochastic decision rule is that demand or cost shocks can result in low prices or profits even when there are no defections. So, the rule to use for determining whether to continue to cooperate has to leave some buffer for uncertainty. The reason for a finite punishment duration is that there is always a positive probability of triggering the punishment, and so the equilibrium should leave room for reversion to cooperative play.

2.4.3 Overlapping Generations of Players

The assumption that play goes on forever is consistent with many institutions, but not with the horizons of the individuals making decisions. However, there are generally older and younger players involved. As long as it is commonly accepted that older players will tend to act more opportunistically and that the younger players will get their turn, then cooperative play is still a SPE. More specifically, a folk-theorem result still applies when (i) there is an infinitely repeated game in which the identities of the players change over time, (ii) each player only considers its payoffs during its term, so that there are no bequest motives, (iii) there is no period in which every player leaves and is replaced by a younger player, and (iv) terms are long, and discounting is slow enough. Should a younger player not tolerate an older player getting a larger share, then the entire play would revert to competitive play for the rest of that player's term. This is enough to support cooperative behavior.

2.5 Summary

This chapter explains how auctions can be modeled as noncooperative games. Single-round or -stage auctions can be modeled as normal form games, and the Nash equilibrium concept is useful in analyzing such auctions. At times, the Nash equilibrium solution will not be able to specify certain outcomes, but rather random ones (mixed strategy equilibria). Such equilibria can be useful in explaining what might otherwise appear to be inconsistent behavior in some auctions.

Extensive form games can be used to represent sequential auctions. Subgame perfect equilibrium will often characterize the plausible outcomes, and

2.5. SUMMARY

ones that rule out noncredible implied threats. However, there can be room for signaling in extensive form games to further refine the equilibrium set. Such signaling will not always be effective, even when there are no costs of conveying signals that can be understood by rival bidders. Costly signaling can, at times be used to sort out different types of outcomes, e.g., those that might be achieved if everyone knew certain bidders were strong, from those that would be achieved if those same bidders were known to be weak.

SPE of repeated games can lead to a wide range of outcomes. Auction rules that limit information available to bidders can make bidding more competitive. However, absent some model of behavior, it can be difficult to predict what type of outcomes are likely in repeated auctions.

$44 CHAPTER \ 2. \ GAME \ THEORY, AUCTION \ DESIGN, AND \ STRATEGY$

Chapter 3

Revenue Equivalence

Summary This chapter explores the relation in outcomes between different auction types. It explains that a first-price sealed-bid auction is strategically equivalent to a Dutch auction. Also, conditions are provided under which an English auction is strategically equivalent to a second-price sealed-bid auction. This chapter also provides an explanation of the derivation of perhaps the single most significant result in auction theory, the revenue equivalence theorem (RET). This result states that the expected revenue and bidder payoffs from two auctions will be the same when the final allocations are the same.

3.1 The Four Basic Auction Types for Single-Object Auctions

A single auction can include one or multiple objects, and can also include multiple units of each object available. Multiple-object auctions, and even multiple-unit auctions for a single product, confront bidders with the problem of anticipating prices in other auctions or for other objects when bidding on any single unit or object. This is explicitly the case in a sequence of auctions for substitutes and/or complements. But it is also the case when there are parallel independent auctions. This chapter focuses on single-object auctions, where these interdependences are not an issue. Much insight can be derived from an analysis of these types of auctions. A great many different auction formats have been used even for single-object auctions. By far the most common are four main traditional types of auctions:

- 1. English, ascending price, open outcry
- 2. Dutch, descending price
- 3. First price sealed bid
- 4. Second price sealed bid

And the English and first-price sealed-bid are by far the two most common auction formats. These auctions are also primarily used for the sale or purchase of a single object at a time. Each of these auctions can be used iteratively to sell multiple objects. This is typically the case at both traditional on-site auctions and in many online auctions. However, as will be seen in subsequent chapters, these four traditional auctions are not well suited for the sale of substitutes or complements, which involve strategic cross-auction interdependences for bidders.

The English auction has several, closely related equivalents. One version has the auctioneer announcing an ascending sequence of prices, and bidders agreeing to keep in the auction or dropping out.¹ A strategy for a bidder in this form of the English auction is a rule to determine when to stop bidding (drop out). In a truly open auction, bidders will see the number of other bidders still remaining, and possibly their identities, when making a decision about whether to drop out. This can be material information and, as is explained in further detail below, can materially affect the outcome. More commonly, auctioneers tend to start the price very high, then quickly backtrack and lower the starting price. In addition, bidders do not have to agree to each announced price (except in the Japanese auction variant), and at times bidders can call out their own price suggestions with larger or smaller increments than is suggested by the auctioneer. Online auctions typically also have the feature of the English auction that bidders are asked

¹This form of auction is sometimes referred to as a Japanese auction, or a Japanese-style English auction.

to respond to an increasing sequence of prices. However, online auctions typically have fixed deadlines for submitting bids.² This makes an online auction strategically much different than a traditional English auction. On the other hand, when the online auction has an extension rule that prolongs the auction by a fixed amount of time after each new bid, the auction is essentially an English auction.

The descending price Dutch auction has the auctioneer start with a high price, and the auction ends when the first bidder enters the auction.

The first- and second-price sealed-bid auctions are auctions in which bidders submit a sealed bid for the object for sale. In both, the high bidder wins. In the first-price auction, the winner pays the bid amount. In contrast, in the second-price auction, the winner pays the highest losing bid amount.

One of the main insights of the early auction literature³ is that the English and the second-price sealed-bid auctions are, in many cases, strategically equivalent, and likewise the Dutch auction and the first-price sealed-bid auction. What is meant by strategically equivalent is that the strategies bidders would want to use are essentially identical.

As is explained below, this strategic equivalence of the English auction and the first-price auction will hold when information revealed during the course of bidding in an English auction will not affect a bidder's valuations or strategies. However, when information a bidder gets during bidding rounds can affect valuations or expectations, then bidding behavior is affected, and the strategy equivalence need not hold.

3.2 Auction Strategy in the Four Basic Auction Types

Each of these four auctions can be described as a normal form game. Sealedbid auctions are very straightforward to describe in this fashion. In both firstand second-price auctions, the bidders (the set of players) have as a strategy

 $^{^2{\}rm This}$ type of auction is sometimes referred to as a Yankee auction. $^3[73]$ and [61] .

the choice of a monetary amount to offer for a single available object. The winner determination in both is the same – the high or best bidder (assuming a forward auction). The only difference is in the payments: in the first-price auction, the winner pays its bid amount. In contrast, in the second-price auction, the winner pays the second highest bid amount. Losers pay nothing.

The Dutch auction is also very straightforward to describe in normal form. In any Dutch auction, the bidders (players) will see a decreasing sequence of prices. The prices may decrease more or less continuously, or in fairly large steps. A strategy is the price at which to first agree to purchase the item. The winning bidder is the one that enters first, and it pays its bid amount. Losers pay nothing.

In an English auction, the strategy can be more complicated. Bidders will face an increasing sequence of prices. A strategy is a decision rule determining when to drop out. The decision rule can depend on the size of the increment and on the number of other bidders remaining in the auction at each price. It can also depend on how fast each rival bidder is increasing its offer. In simple auctions, when prices increase continuously or nearly so, then bidders will base the decision to stop bidding on their own values, and not on the size of the increment. If the outcome of other bidders is immaterial or unobservable, then the strategy for when to stop bidding is simply a limit price. In more complex auctions in which rivals' information or identities are material, the bidding strategy can depend on who is still bidding in the auction. The winner is the last bidder in, and at the last agreed-on price.

The Japanese auction variant of the English auction requires all bidders to remain in at each posted price as long as they are willing to purchase the item. Once a bidder drops out, it is no longer able to reenter the bidding. As in the standard English auction, the activity of other bidders can materially affect any given bidder. But, if not, the strategy in the Japanese auction is simply a stopping price.

3.3 Strategic Equivalence

This section explains the strategic equivalence of first-price auctions and Dutch auctions, and of second-price auctions and English auctions.

3.3. STRATEGIC EQUIVALENCE

First consider the English auction. As noted above, in the most usual English auction, an auctioneer calls out an escalating sequence of prices, and at least one bidder that is not the standing high bidder must meet the most recently announced price, within a fixed amount of time, to keep the auction from ending. In this type of auction, the optimal strategy for a bidder with a value v is to keep bidding as long as the price is no higher than v, and then drop out.⁴ This assumes that bid increments are small; if not, then a bidder may have an incentive to skip a turn responding to a new bid. This also assumes that a bidder does not revise its valuation during the auction based on what it sees from others' behavior. This is not always the case.

For example, in a common-value auction, each bidder has a signal of the true value $s_j = v + \epsilon_j$, where v is the true value and ϵ_j is an independently, identically distributed error term. This type of situation is common in oil lease bidding.⁵ In this case information about others' bids provides some information about a bidder's forecast error. In the case of oil lease bids, it was observed that although bidders may have, on average, unbiased forecasts, they will only win when their forecasts are too optimistic. Seeing other bidders staying in or dropping out provides information about how optimistic is a given bidder's forecast; this information can be used to revise expectations. This type of revision is not needed in an auction in which bidders have independent private values, that is, each bidder's value is uncorrelated with its rivals'.

In the second-price sealed-bid auction, it is straightforward to see that a bidder's dominant strategy is to bid its true value.⁶ There is no incentive to bid more than one's value. For, suppose a bidder offers b > v, and the highest offer by a rival is b'. There are three cases to consider:

1. $b' \leq v < b$. Then overbidding does not matter. This bidder wins and pays b'.

 $^{^4\}mathrm{How}$ quickly a bidder should bid can also be part of the strategy, as it can influence other bidders.

 $^{^{5}}$ See [17].

⁶The following discussion assumes independent private values. If there are affiliated values, it is still a dominant strategy for a bidder to bid its value, but this value is conditional on it being the high bidder.

- 2. $v < b' \leq b$. Then this bidder wins, and pays b' > v. In this case, overbidding reduces this bidder's payoff.
- 3. b' > b. Again overbidding does not matter, as this bidder loses.

A similar argument can be used to show a bidder will never want to bid any b < v. If it did, there is a chance it would lose, when it should not. On the other hand, if bidding b' > b does not affect whether the bidder wins, then this higher bid would impose no additional costs on the bidder. Thus, underbidding cannot be an equilibrium. This shows that in a second-price sealed-bid auction, it is always optimal to bid one's value; in other words, this is a dominant strategy.

What this means is that the point at which a bidder stops bidding is the same in an English auction as the amount a bidder would offer in a secondprice sealed-bid auction. Further, the bidder pays the second highest bid amount. Thus, the auctions are strategically equivalent, and achieve the same prices and payoffs to the bidders. The auctions also give the seller the same price, and the same buyer pays it. Thus these two auctions are also *outcome* equivalent, as well as *payoff* equivalent, at least in an independent private values setting. Note that if bidders cannot observe rivals' bids, and price increases are continuous, the fact that the auction continues can convey information. When bidders are never told who else is bidding, or how many other bidders remain active, and the stopping rule is for the auctioneer to continue increasing price until a fixed upper bound higher than market clearing, then the English and second-price auctions will always be equivalent – strategically and with respect to payoff and outcome.

In a first-price sealed-bid auction a bidder will want to bid less than its value. If a bidder offers its full value, then it will earn zero surplus even when it wins. The calculation of how much a bidder should offer involves a tradeoff – a lower offer provides a higher surplus if the bidder wins, but a lower probability of winning. The bidder will want to to choose its bid amount, b, so as to maximize

$$P(b)(v-b) \tag{3.1}$$

where v is the bidder's value and P(b) is the probability that it wins.⁷

In a Dutch auction, a bidder will not want to bid until the price falls at least to its value v. Once the price falls to v, the bidder will want to choose a price to make a first offer b so as to maximize (v-b)P(b). This is the same expression as (3.1). Thus, a bidder will want to choose the same bid amount in both first-price sealed-bid auctions and Dutch auctions. The Dutch and the first-price sealed-bid auctions are strategically equivalent: bidders will have incentives to bid the same amounts.

In summary:

Theorem 1 ([61] and [73]. In (reduced) normal form, Dutch auction and first-price sealed bid auction are strategically equivalent.

A strategy in the first-price sealed-bid auction is just the price the bidder offers. In a Dutch auction, the strategy is the price at which to start bidding. Assuming that price falls continuously in the Dutch auction, the expression (3.1) establishes that these are the same decisions.

Revenue Equivalence of English and Dutch 3.4Auctions

One question often raised is: What difference does auction design make? For instance, there has been much, and sometimes heated, debate about the relative merits of uniform price vs pay-as-bid auctions in the energy sector.⁸ This section explains one of the more remarkable theoretical results of economics – the revenue equivalence theorem (RET). The result is that the expected payments and outcome of the English and Dutch, and therefore those of the first-price and second-price auctions, are the same – at least when bidders have independent private values.

⁷The first-order conditions are that the optimal b will satisfy (v - b)P'(b) - P(b) = 0, or $\frac{v-b}{8} = P(b)/bP'(b)$. ⁸See [40].

Recall that in an English auction, or a second-price sealed-bid auction, all bidders bid truthfully; in equilibrium, the high-value bidder wins and pays (approximately) the valuation of the bidder with the second highest valuation. In the Dutch auction, or the first-price sealed-bid auction, the high bidder will still usually win,⁹ but the amount of the winning bid, which is the transaction price, will generally be less than the value. The question is by how much the winning bid amount will fall short of the winning bidder's value. The remarkable answer is that this gap is exactly equal to the expected value of the second highest value of the bidders, or the equilibrium price in an English auction – at least in the case in which bidder values are independently distributed.

The intuition for this result is fairly straightforward, although the derivation is somewhat more complicated than for strategic equivalence of the bid strategies in the Dutch and first-price, sealed-bid auctions. Consider the first-price sealed-bid auction. The high bidder won't know ex ante who the high bidder is. However, this high bidder should only be concerned with the case in which it is the high bidder, as in other cases its bid won't matter.

So, the high bidder will want to bid just a bit more than the second highest bid amount. This amount is no more than the second highest value, as the second highest bid can never be above value in equilibrium in a first-price sealed-bid auction. Of course, the second highest bidder will offer a bit less than its value too. So, the high bidder will want to weigh the benefit of bidding a bit more than the expected amount of the second highest bidder's value. This high bidder would gain a bit of surplus by bidding a bit less, but would lose with a higher probability. When the high bidder's offer exactly equals the expected amount of the second highest bidder's value, then the benefits of increasing its bid a small amount will exactly equal the costs.

This RET theorem can be established in two different ways. The more direct approach is to compute optimal bid functions, and then to show that the optimal bid in equilibrium means that the high bidder will bid the expected amount of the second highest bidder's value. The next section develops an approach to calculating an optimal bid function, but does not provide a complete proof the RET.

 $^{^{9}\}mathrm{It}$ is assumed that bidders use equilibrium bid strategies, and have the same preferences toward risk.

The less direct, and much more general approach is to rely on what is called the *revelation principle*, to restrict attention to *direct* revelation mechanisms, that is, auctions which essentially ask bidders to report their type. In direct revelation mechanisms, the bidders' reports determine the winner and the payment. Auctions are a specific form of direct revelation mechanism – bidders are asked to report values, and payments are based on those reports. Of course in a first-price sealed-bid auction a bidder won't wish to offer its full value, but its offer can be used to impute its value. In that sense a first-price sealed-bid auction is a direct revelation mechanism.

The RET states that any two auctions which always have the same winners, and give the lowest-value bidder the same (zero) expected surplus, always result in the same auction revenues. The intuition for this result is as follows.

Consider a bidder with value v, who submits a bid $b = \beta(v)$, where $\beta(v)$ is a bid function. Now, consider the incentives of this bidder to report that it is a slightly stronger or weaker bidder by reporting $b + \Delta = \beta(v + \delta)$. It has to be the case, at equilibrium, that this bidder will not gain by reporting a bit more or less. So, the extra payoff is offset by the extra probability of losing by altering a bid. This would have to be true for any two auctions that always result in the same winners. This means that the shape of the total payment schedule, P(v), would have to be same for any two auctions that always result in the same winners. And if the lowest-value bidders always get a zero payoff, then the two auctions would always generate the same revenues. [41] provides a more general statement.

Theorem 2 ([41] p. 43). Assume each of N risk-neutral potential buyers has a privately known value independently drawn from a common distribution F(v), which is strictly increasing and atomistic on $[\underline{v}, \overline{v}]$. Also, suppose no buyer wants more than one of the k available and indivisible objects. Then any auction mechanism in which (i) the objects always go to the k buyers with the highest values and (ii) any bidder with value at \underline{v} expects zero surplus yields the same expected auction revenue and results in a buyer with value v making the same expected payment.

Sketch of Proof

Suppose values are iid draws on $[\underline{v}, \overline{v}]$. Let F(v) be the distribution, with $F(\underline{v}) = 0$ and $F(\overline{v}) = 1$.

Let $S_j(v)$ be the expected surplus of bidder j with value v. Let $P_j(v)$ be the probability that bidder j wins.

So $S_j(v) = vPj(v) - E$ [payment by type v of player j]. For, if bidder j reports any $v' \neq v$, then $S_j(v') = P_j(v')v$ – [expected payment by a bidder j of type v'].

So
$$S_j(v) \ge S_j(v') + (v - v')P_j(v')$$
.

This implies $S_j(v + dv) \ge S_j(v) + (-dv)P_j(v + dv)$. Also, $S_j(v + dv) \ge S_j(v) + dvP_j(v)$.

The above two inequalities imply $P_j(v+dv) \ge [S_j(v+dv)-S_j(v)]/dv \ge P_j(v)$. Letting $dv \to 0$ leads to $\frac{dS_j}{dv} = P_j(v)$. (Note that if S(v) = vP(v) - X(v), where X(v) is the payment that the bidder must make, the condition becomes $\frac{d[v_jP(v)-x(v)]}{dv} = P_j(v)$.)

Integrating up to v yields

$$S_j(v) = S_j(\underline{v}) + \int_{\underline{v}}^{v} P_j(x) dx$$

Now, consider any two mechanisms that have the same $S_j(v)$ and the same $P_j(v)$ for all v and for all players j. So, j's expected payoff is the same for both mechanisms. This means that the auctioneer gets the same expected revenues. \Box

A direct application of the above result is the revenue equivalence between first-price, second-price, English, and Dutch auctions when bidder valuations are independent draws from the same distribution. This result is also key to characterizing an optimal auction, as is explained in more detail in the next chapter. Absent any reserve price, the second-price or English auction yields an expected revenue of the the value of the second highest valuation among the bidders. So, at least in the case of independent private values, a second-price auction with a positive reserve price will be optimal.

3.5. SUMMARY

The above statement of the RET imposes a few significant assumptions. One is that each bidder is only interested in at most one unit. Should a bidder want more than one unit, strategic withholding, that is, bidding low on one unit hoping to affect the clearing price can be a factor. This possibility is discussed in more detail in Chapter 8. Another key assumption is that buyers have independent private values. As is explained in Chapter 5, RET no longer holds in a *common value* auction, that is, an auction in which values of different buyers are correlated.

3.5 Summary

This section considered single-object auctions. The main insight, that the auction rules for determining the winner and the price paid can affect bidder incentives, was used to characterize some key properties of auctions. First, there is a strategic equivalence between Dutch and first-price sealed-bid auctions. Second, under some assumptions, there is also a strategic equivalence between standard English auctions and second-price sealed bids. The more significant result is that the auction revenues can be the same for two auctions that appear quite different, such as the English auction, the second-price auction, the Dutch auction, and a standard first-price sealed-bid tender.

The RET is significant in practice. In particular, it suggests that discriminatoryprice auctions, in which each winner pay its own bid amount rather than same price others pay for identical lots, will not always produce better revenues or results than uniform-price auctions. As noted above, this means that decisions about the payment rule that occur in some sectors, such as energy, should focus on whether revenue equivalence holds and on bidder incentives, and not on the fact that bidders pay less than bid amounts in a uniform-price auction. Also, spectrum auctions that provide some bidders with discounts can create strong incentives for bidders to find ways to qualify, with the net effect of bidding away such discounts.

This has been the experience in the US FCC spectrum auctions, where bidding credits have been competed away. The experience in FCC Auction 11 illustrates this quite clearly (Figure 3.1). That auction included three identical licenses in each of 493 geographic regions, called Basic Trading Areas.

	Market		High	High Bid	Net Bid	Price	Ave Price D/E -	
Bloc	k Name	Num Bids	Bidder	(\$k)	(\$k)	(\$/bu)	Price of F	% difference
E	New York, I	0	AT&TWire	58,800	58,800	0.33		
F	New York, I	0	NorthCoast	100,320	75,240	0.56		
D	New York, I	0	OPCSE	50,700	50,700	0.28	(\$20,490,001)	-40%
F	Los Angele:	0	AerForce	5,965	4,474	0.04		
D	Los Angele:	0	AT&TWire	37,510	37,510	0.26		
E	Los Angele:	0	Rivgam	31,910	31,910	0.22	\$28,745,000	90%
F	Chicago, IL	0	NextWave	30,753	23,065	0.38		
D	Chicago, IL	0	SprintCom	59,976	59,976	0.73		
E	Chicago, IL	0	SprintCom	62,741	62,741	0.77	\$30,605,499	49%
D	San Francis	0	AT&TWire	13,655	13,655	0.21		
F	San Francis	0	NextWave	5,779	4,334	0.09		
E	San Francis	0	Western	10,737	10,737	0.17	\$7,861,753	73%
D	Philadelphia	0	Comcast	12,169	12,169	0.21		
F	Philadelphia	0	NextWave	29,407	22,055	0.50		
E	Philadelphia	0	Rivgam	12,761	12,761	0.22	(\$9,590,250)	-75%
D	Detroit, MI	0	NextWave	3,815	3,815	0.08		
E	Detroit, MI	0	OPCSE	3,856	3,856	0.08		
F	Detroit, MI	0	OPCSE	8,500	6,375	0.18	(\$2,363,081)	-28%
D	Dallas, TX	0	AT&TWire	25,895	25,895	0.60		
E	Dallas, TX	0	AT&TWire	27,060	27,060	0.62		
F	Dallas, TX	0	NextWave	21,340	16,005	0.49	\$1,695,000	8%
F	Boston, MA	0	NorthCoast	8,909	6,682	0.22		
D	Boston, MA	0	OPCSE	6,515	6,515	0.16		
E	Boston, MA	0	OPCSE	7,515	7,515	0.18	(\$1,893,996)	-25%
F	Washingtor	0	AerForce	11,780	8,835	0.29		
E	Washingtor	0	OPCSE	6,071	6,071	0.15		
D	Washingtor	0	Rivgam	6,820	6,820	0.17	(\$5,334,493)	-78%
E	Houston, T	0	AT&TWire	9,835	9,835	0.24		
D	Houston, T2	0	SprintCom	13,259	13,259	0.33		
F	Houston, T2	0	Telecorp	10,150	7,613	0.25	(\$1,869,539)	-18%

Figure 3.1: FCC Auction 11

One of the three licenses, the F block, was reserved for "designated entities" (small bidders), and the other two blocks were available to all bidders. The small bidders also received bidding credit, that is, they paid a discounted price. The net effect was that the small bidders paid virtually the same prices as the larger ones.¹⁰

 $^{^{10}\}mathrm{Figure}$ 3.1 shows the results of the bidding in the largest ten regions.

Chapter 4

Optimal and VCG Auctions

Summary

This chapter provides a description of how the optimal bid strategy is calculated for a bidder. This is then applied to characterize optimal auction design. Finally, this chapter provides a characterization of all auctions in which bidders will have incentives to accurately report true valuations, that is, Vickrey-Clark-Groves auctions.

4.1 Optimal Auctions

This chapter characterizes an optimal auction. "Optimal" here is used to mean an auction design and configuration that, for a forward auction, maximizes expected revenues. It turns out that most often many designs will achieve the same expected outcome. This should be apparent from the revenue equivalence theorem, which states that the English, Dutch, first-price, and second-price auctions all result in the same outcome in the case in which all bidders have independent private values. The reservation price is a configuration parameter of all four main auction formats discussed in the previous chapter, and it can affect the outcome: the level of the reservation price can determine whether or not an object is sold. On the other hand, a public and secret reservation, under the assumption of independent private values, should result in the same expected revenues, assuming that all bidders will participate and higher-value bidders always offer more than lower-value bidders. In practice, bidders face participation costs, and so a private reservation which reduces participation can influence the outcome and auction revenues. This is also not to suggest that any two auctions will always result in the same expected revenues, even assuming the same winners and the same reservation prices. As will be explained in Chapter 5, when bidders values are affiliated, the English auction generates higher revenues than does either a first-price or a second-price sealed-bid auction.

4.2 Calculating Optimal Bid Functions

This section shows how to directly calculate optimal bid functions for a bidder seeking to maximize its expected surplus in a first-price auction. This calculation can take two approaches. The first is to rely on RET, and estimate the amount of the second highest bidder's value. This approach works for Dutch auctions and first-price auctions when bidders have independent private values. The second, more general approach is to directly optimize the expected bidder surplus.

The following describes how optimal bid strategies can be calculated for the case in which there are N symmetric bidders competing in a first-price auction. Each bidder j's valuation, v_j , is, ex ante, random. The density function of the realized valuations is the same for all bidders, is denoted f(v), and the corresponding distribution is F(v). A strategy is a function $\beta(v)$ mapping a bidder value into bids. Then, in this symmetric case, each bidder of type v will choose a bid b to maximize $(v - b)F^{N-1}[\beta^{-1}(b)]$.

This implies the first-order conditions

$$-F^{N-1}\left(\beta^{-1}(b)\right) + (N-1)(v-b)f\left(\beta^{-1}(b)\right)F^{N-2}\left(\beta^{-1}(b)\right)\beta'^{-1}(b) = 0$$

As $\beta'(\beta^{-1}(b)) = \frac{1}{\beta'(v)}$, where $b = \beta(v)$, the first-order condition can be rewritten as

$$-F^{N-1}(v) + (N-1)(v-b)F^{N-2}(v)f(v)/\beta'(v) = 0$$

or

$$\frac{d}{dv}\Big[(v-\beta(v))F^{N-1}(v)\Big] = F^{N-1}(v)$$

Integration implies

$$\beta(v) = v - \frac{K + \left\{ \int_0^v F^{N-1}(s) ds \right\}}{F^{N-1}(v)}$$
(4.1)

Note that the condition that a zero-value bidder gets zero implies K = 0.

It is relatively straightforward to calculate optimal bids when values are uniformly distributed. Let [v, V] denote the interval over which values are distributed, and L = V - v. Then the expected value of the high bidder's valuation is $\frac{N}{N+1}L + v$, and this bidder's expected price is $\frac{N-1}{N+1}L + v$, which can be derived from (4.1).

The simplest case is when there are two bidders whose have values uniformly distributed over some interval, say [0, M]. In this case, a bidder with value v will wish to choose its bid, b, to maximize $(v - b)\frac{b}{v}$. This leads to a solution $b = \frac{v}{2}$.

To calculate optimal bid functions more generally using this direct approach requires an explicit formula for the distribution function F(v). The indirect, RET approach offers a shortcut, through calculating the amount of the second highest bidder's valuation. An exact calculation of this latter amount may be no easier than the direct approach, but to the extent that one can approximately estimate the value for the second highest bidder, that will make it easier to determine how to bid.

These approaches also assume that bidders are risk neutral. Consider a bidder with value v, who can offer either b or $\tilde{b} > b$. Let ρ be the probability of winning with an offer of b, and $\tilde{\rho} > \rho$ be the probability of winning with an offer of \tilde{b} . The expected payoff from bid b is is $E(b) = \rho(v-b)$, and from bid \tilde{b} is $E(\tilde{b}) = \tilde{\rho}(v-\tilde{v})$. In words, the offer of \tilde{b} results in a lower payoff, with a

difference $\tilde{b} - b$, but a higher probability of winning, $\tilde{\rho} - \rho$. This bidder can strictly prefer \tilde{b} to b, even when $E(b) = E(\tilde{b})$, or vice versa, depending on the bidder's preferences toward risk. If the bidder is more concerned about winning than about the expected payment, it would prefer to bid more than maximizes E(b).

4.3 **Optimal Auctions**

This section summarizes the key results from Myerson's famous (1981) paper on optimal auction design. An optimal auction, that is, one that maximizes expected revenues for the seller, can be derived using two key insights from Myerson's paper. The first insight is that for any auction there is an equivalent auction in which bidders are asked to report values. The second is that the optimal bid strategy determines the shape of the function for the payoff based on reported type, up to an additive constant. The constant can be set to make the lowest-value bidder indifferent between participating and not participating in the auction. This fully defines the optimal auction.

4.3.1 The Revelation Principle

Auctions for a single object can take many forms, and both losers and winners can be required to make payments. The following assumes, as above, that each bidder *i*'s type, or value, v_i , is distributed according to a probability density function, $f(v_i)$. The values are distributed over some range $[a_i, b_i]$ for i = 1, 2, ..., N. A strategy for a bidder then is a choice of some action, from a set of permitted actions defined by the rules, i.e., $\alpha_i = \alpha_i(v_i) \in A_i$.

An auction "mechanism" is defined by the strategies and the rules for determining winners and payoffs. A strategy, in the context of an auction mechanism, is a choice of a bid amount as a function of the bidder's type, v. The winners are generally the highest value bidder, or bidders - in a multiobject auction. And a bidder's *payoff* is the value of winning (assuming it wins) less the payment it has to make. A direct revelation mechanism is an auction in which bidders report values; the probability of a bidder winning, $\rho(\mathbf{v})$, and its payment, $x(\mathbf{v})$, will be a function of the vector of reported types of all the bidders, \mathbf{v} . Note that the actual type can differ from the reported type.

The revelation principle states that given any feasible auction mechanism, there is an equivalent direct revelation mechanism which gives the seller and all bidders the same expected utilities as the given mechanism. The reason that the revelation principle holds is straightforward. Suppose bidder j of type v_j in an auction would have chosen action $a(v_j)$. Then, an allocation rule that gives each bidder the probability of winning $p_j =$ $p_j[a_1(v_1), a_2(v_2), \ldots, a_n)v_n)]$ and allocation of $x_j = x_j[a_1(v_1), a_2(v_2), \ldots, a_n)v_n)]$ would provide each bidder the same incentives and outcomes as the given auction. In other words, one can back out strategies bidders would choose to determine prices and allocations consistent with those that would result from bidders' optimal choices in the original auction.

4.3.2 The Revelation Principle and Optimal Auctions

The revelation principle was derived as a lemma in Myerson's 1981 paper and used in solving for optimal auctions. The payment rule and the rule determining the probability a bidder of a given type wins determine the bidders' behavior and the auction's outcome. This follows because the payment rule and the probability of winning determine bidding incentives. A bidder will prefer to report being of type v to any other type $v + \Delta v$ when $S(v) \geq S(v + \Delta v) - \Delta v P[(v + \Delta v)]$, where S(v) is the net surplus a bidder of type v receives if it wins. As shown above, this condition fully defines the surplus function, and therefore the payment rule, up to an additive constant, $S_i(\underline{v})$, which in turn determines the payment of the lowest-value bidder:

$$S_j(v) = S_j(\underline{v}) + \int_{\underline{v}}^{v} P_j(x) dx$$

Now, consider any two mechanisms that have the same $S_j(v)$ and the same $P_i(v)$ for all v and for all players j. So, j's expected payoff is the same for

both mechanisms. This means that the auctioneer gets the same expected revenues.

Therefore, the term $S_j(\underline{v})$ determines the expected payoff to the lowest-value bidder that actually bids. Variations in $S_j(\underline{v})$, through bidder participation fees, reservation prices, etc., will affect the total auction revenues. Thus, the auction revenues can be written as a function of $R = R(S_j(\underline{v}))$. The optimal auction, that is, the one that maximizes seller's expected revenues, corresponds to the choice of an expected surplus for the lowest-value bidder. This, in turn, is equivalent to the choice of a reservation price.

To see what this value is, consider the case where the seller does not have any value for keeping the object, and where the bidders have values uniformly distributed on [a, b]. An optimal auction will be a modified Vickrey auction in which the seller sets a reservation price.

This is no longer a Vickrey auction, as the seller will have a preference to set a minimum price for the case in which there is only one high bidder. The simplest case is that of one bidder – the seller will want to choose a reservation price of $\frac{b-a}{2}$. This means that the seller will sell the object with probability 0.5. Another way to see this is to look at expected revenue as a function of the reservation price r. This will be r[1 - F(r)], and it is maximized at $r = \frac{b-a}{2}$.

The following calculations illustrate why the auctioneer will want to set a positive reservation price and risk the object not being sold. Suppose the values are uniformly distributed on [0, 1]. When there is no reservation price and one bidder, the equilibrium in the totally noncompetitive "auction" is for this bidder to bid a bit more than 0 in a first-price auction, and its value in a second-price auction. This provides the auctioneer with zero expected revenues. Now, if the auctioneer sets a reservation price, its expected revenues are the reservation price times the probability that this bidder has a value above that level. The optimal reservation price is $\frac{1}{2}$.

When there are two bidders whose valuations are uniform draws from the interval [0, 1] and no reservation price, the expected revenue in a second-price auction is $\frac{1}{3}$. On the other hand, if the auctioneer sets a reservation price of $\frac{1}{2}$, it will receive an expected revenue of $\frac{5}{12}$.¹

¹ To see this, notice that both bidders will have valuations above $\frac{1}{2}$ with probability $\frac{1}{4}$.

More generally, for any n the seller will want to choose this same reservation price at $\frac{b-a}{2}$ when bidders have uniform values on an interval [a, b]. The probability of the object not being sold is $\frac{1}{2}^n$. For a large number of bidders this probability gets arbitrarily small. When the distribution of values for different bidders can have different upper and lower bounds, it is possible that the revenue-maximizing auction will discriminate by allowing bidderspecific reservation prices, so that a bidder who has a higher value, when ex ante that bidder *is* expected to have a highe value, may still lose to a bidder with a lower value.

Also, notice that an auction that maximizes expected revenues will not be efficient. When an object has no intrinsic value for the seller, then the efficient outcome is for the seller to award the object to the bidder with the highest value for it. An open first- or second-price auction with no reservation will be efficient in this sense. However, this will not maximize the seller's expected revenues. The seller will want to impose a positive reservation price.

When a government agency is auctioning off a resource, or concession, such as a spectrum license, oil lease, or highway concession, the optimal auction will not usually be the one that maximizes revenues. The agency may have other objectives, such as ensuring that the object is acquired and used. For instance, if an agency wants to promote development of the communications sector, it is likely to prefer allocating spectrum for free, if no one is willing to pay a positive price, with some sort of buildout requirement. This suggests that a government agency should set a zero reservation – unless the revenues lost from the optimal auction would need to be made up from a different source, such as increases in tax rates. Indeed, there are times when an agency will set a negative reservation price; for example, in allocating spectrum or telecommunications concessions agencies may want to provide subsidies to ensure target coverage goals are met.

This analysis assumes all bidders show up. If bidder participation can be random or if there are participation costs, then another rationale for setting a positive reservation is that the auctioneer would have the possibility of withdrawing an object from the auction and reauctioning at a later date.

In this case, expected revenues are $\frac{2}{3}$. They both have valuations below $\frac{1}{2}$ with probability $\frac{1}{4}$, resulting in the object being unsold. And then, with probability $\frac{1}{2}$, only one bidder can meet the reservation price, and revenues equal the reservation price. This results in expected revenues of $\frac{1}{4} \times \frac{2}{3} + \frac{1}{2} \times \frac{1}{2} + \frac{1}{4} \times 0 = \frac{5}{12}$.

When participation at each auction date is random, then, as explained in more detail below in the analysis of multi-unit auctions, a positive reservation price could increase revenues and efficiency. This is very clearly the case when the auctions are conducted for future delivery of an object.

4.4 Incentive-Compatible Auctions

This section describes auctions which provide each bidder an incentive to report true values. The simplest form of such an auction mechanism is the second-price auction for a single object. For single-object auctions, the second-price sealed-bid auction is a mechanism under which it is a dominant strategy for bidders to report true values, that is, the auction is *incentive compatible*. What is key to this incentive compatibility is that a bidder's increasing its offer has no affect on the price the bidder pays *except* when this increase affects whether the bidder would have won, that is, when the bidder is *pivotal*. What is shown below is that this idea can be extended to multi-object auctions.

4.4.1 The Vickrey-Clark-Groves Mechanism

Recall that, from the RET, a second-price and a first-price auction for a single object will, with independent private values, result in the same expected outcome. However, it is *not* the case that bidder behavior will be the same. Vickrey noted that in a second-price auction each bidder has a dominant strategy to truthfully report the object's value. In contrast, in a first-price auction, each bidder will pay its bid, and each bidder will offer an amount less than its value. In equilibrium, the amount by which the winning bidder will shade its bid is the expectation of the second highest value. This section shows is that in more general settings with multiple objects it is possible to fully characterize all auctions in which bidders will always want to bid truthfully. Such auctions are called Vickrey-Clark-Groves (VCG) mechanisms. A second-price auction is just a special case of a VCG mechanism.

What follows assumes *additively separable* utilities, that is, a bidder's utility of winning an object or allocation x can be written as $u_i = v_i(x) + t_i$, where

 t_i is the monetary transfer that *i* receives. The revelation principle states that it suffices to consider only direct revelation mechanisms in characterizing auctions in which bidders will have incentives to truthfully report valuations. In other words, bids can be assumed to take the form of reported valuations for an outcome or an allocation.²

A VCG auction is a mechanism in which a bidder's transfer satisfies the following condition:

$$t_i[\mathbf{w}(\mathbf{x})] = \sum \mathbf{w}_{-i}[\mathbf{x}^*(\mathbf{w})] + \mathbf{h}_i(\mathbf{w}_{-i}(\mathbf{x}))$$
(4.2)

where $x^*(\mathbf{w})$ is the allocation when $\mathbf{w}(\cdot) = (w_1(\cdot), \ldots, w_n(\cdot))$ is the vector of the reported valuations and \mathbf{w}_{-i} is the vector of valuations of all bidders other than i.

Note that i's bid only affects its transfer, or payment, in (4.2) to the extent that i's bid changes the allocation that other bidders receive. If i's bid is not pivotal, then i's bid has no effect on its monetary payment. Further, it can be shown that a mechanism that satisfies (4.2) has the property that i's transfer will equal the amount by which i's bid affects the aggregate welfare of the other bidders, i.e.,

$$t_i(\mathbf{w}_{-i}, w_i) - t_i(\mathbf{w}_{-i}, w'_i) = \sum \mathbf{w}_{-i}(\mathbf{x}^*) - \sum \mathbf{w}_{-i}(\mathbf{x}')$$
(4.3)

where x^* is the allocation at (w_{-i}, w_i) and x' is the allocation when the reported valuations are (w_{-i}, w'_i) . In other words, *i*'s report only affects *i*'s transfer to the extent that this report affects the aggregate welfare of other bidders. Indeed, this latter condition is essentially a necessary and sufficient condition for an auction mechanism to provide bidders an incentive to bid true values.³

²In Chapter 10, properties of VCG auctions are explained in further detail.

³See [31] for a more formal discussion.

4.4.2 Properties of VCG Auctions

The second-price auction and, more generally, the VCG auction has the desirable property of eliciting truthful revelation of preferences and of achieving efficiency. This raises the question why the VCG auction should not be used more generally. This section addresses this issue.

One property of VCG auctions is that a winning bidder does not have to pay its bid amount. The gap between the winning bid and the payment can be quite large. Auctions in which the second price was less than 2% of the winning bid occurred among the first spectrum auctions.⁴ This gap does not mean that there was necessarily any lost revenue. However, it does suggest that a reservation price or a higher reservation price might have been advisable. Such large gaps can look bad for the auctioneer and also require post-auction rationalization, they invite challenges and litigation, and they may result in cancellation or re-auction. A first-price auction or an English auction avoids this second guessing. And, in some situation, an English auction will also result in higher revenues.

VCG and Budget or Revenue Constraints

The VCG auction has a number of other properties that may discourage its adoption. Recall that in a VCG auction, a bidder payment is given by (4.2). One condition that might be needed to ensure feasibility of the auction is some budget balance requirement. Or, in a government-managed auction, the revenues might be capped or required to be distributed back to the participants net of administrative costs. For example, the auction can include contributions by users and offers from suppliers. The agency conducting this (two-sided) auction might not be allowed to incur a deficit. Absent some additional restrictions, the sum of the bidder payments in (4.2) can otherwise be arbitrary. This limitation on transfers to bidders imposes the condition that the sum of the terms in (4.2) is zero. However, as [54] has shown, this type of restriction can be incompatible with the requirement for a VCG auction to exist in (4.3).

⁴See Mueller [56].

To see this, consider the case of two bidders for one object.⁵ The required condition is that there are functions h_1 and h_2 that always add up to zero Suppose that bidder 1's possible values for the object are either 1 or 3, and bidder 2's possible values are either 2 or 4.

- If bidder 1's valuation is 1 and bidder 2's valuation is 2, then bidder 1 should pay $p_1 = h_1(2)$ and bidder 2 should pay $p_2 = 1 + h_2(1)$. So we seek $p_1 + p_2 = 0 = 1 + h_2(1) + h_1(2)$.
- If bidder 1's valuation is 3 and bidder 2's valuation is 4, then it must be the case that $p_1 + p_2 = 0 = 3 + h_2(3) + h_1(4)$.
- If bidder 1's valuation is 1 and bidder 2's valuation is 4, then it must be the case that $p_1 + p_2 = 0 = 1 + h_2(1) + h_1(4)$.
- If bidder 1's valuation is 3 and bidder 2's valuation is 2, then it must be the case that $p_1 + p_2 = 0 = 2 + h_2(3) + h_1(2)$.

Adding of the first two and last two equations from the above list requires that

$$0 = 4 + h_2(1) + h_1(2) + h_2(3) + h_1(4) = 3 + h_2(1) + h_1(2) + h_2(3) + h_1(4)$$

This is impossible, so budget balance fails. More generally, this means that a VCG auction may fail to exist when the auction must also meet specific revenue targets, even when the revenue targets are consistent with bidder valuations, such as raising zero net revenues in the above discussion.

4.4.3 Multi-object VCG Auctions

This section provides a brief description of some multi-product VCG auctions; Chapter 10 provides more detailed analysis of multi-object VCG auctions. Where bidders are each seeking to purchase at most one object, the second-price auction can be readily modified for the case of k identical objects. More specifically, when all objects are sold to the k highest bidders,

⁵The following is from Milgrom [54].

but for the k + 1st highest price, then each bidder will have as a dominant strategy to bid its true value, as is the case in the standard second-price auction.

Spectrum Auction Examples

Several spectrum auctions have been conducted using a modified VCG format.⁶ There have been a few auctions of identical spectrum licenses which use a variation of the k + 1st-price rule, but are no longer incentive compatible. One was in Denmark in 2011.⁷ That auction included four generic licenses, and bidders could win one license each. The four highest bidders won, but paid the fourth highest bid amount. In this case it was no longer true that a bidder would always want to bid its value. When it is possible for a bidder to be a marginal bidder, then the bidder would want to shade its offer, just as in a standard first-price auction.⁸

Lot	Winner	High Bid	2nd Bid
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 A.B.	\$401,000	\$100,000
7	United Christian	\$685,200	\$401,000

Table 4.1: New Zealand 8 MHz UHF TV license auction

Another series of spectrum auctions in New Zealand used a second-price rule (Mueller ([56])). However, these auctions were not incentive compatible, as bids for each identical license were taken separately, and the price for each license was the second highest bid for that license. Table 4.1 from [56] shows

 $^{^6}$ See http://stakeholders.ofcom.org.uk/spectrum/spectrum-awards/ for examples of such auctions in the UK.

⁷See http://en.itst.dk/spectrum-equipment/Auctions-and-calls-for-tenders/3g-hovedmappe/3g-auction-2001-1

⁸There was no document released by the Danish government indicating the reason for this provision. It may be the case that this auction design could result in higher prices, or, at least, it would be less obvious when a bidder paid much less than its value.

4.4. INCENTIVE-COMPATIBLE AUCTIONS

one such *simultaneous* auction for seven identical licenses. The prices paid ranged from NZ\$100,000 to NZ\$401,000.

The New Zealand government also conducted a set of *separate* second-price auctions for three cellular licenses. The relevant bids for these three auctions were as follows: The following tables - 4.2, 4.3 and 4.4 show the results.

Bidder	Bid amount
Telecom New Zealand	NZ\$101,200,000
First City Capital	NZ\$11,158,800
Imagineering Telecommunications	NZ\$1,388,000

Table 4.2: New Zealand AMPS A Auction

Bidder	Bid amount
Bell South	NZ\$85,552,101
Telecom New Zealand	NZ\$25,200,000
Racal-Vodafone Ltd	NZ\$1,000,000
Broadcast Communications Ltd.	NZ \$2,000

Table 4.3: New Zealand TACS A Auction

Bidder	Bid amount
Bell South*	NZ\$85,552,101
OTC International [*]	NZ\$13,250,000
Telecom New Zealand [*]	NZ\$7,000,000
Broadcast Communications Ltd.	NZ\$5,000
Michael Oliver Thaisen	NZ\$300

Table 4.4: New Zealand TACS B Auction

The winning bidders had to pay NZ\$11,158,800 and NZ\$25,200,000 for the AMPS-A and TACS-A licenses. The three high bidders for the TACS-B license (marked *) were disqualified, and the entire auction was rerun. In these auctions for cellular licenses, the bidders had to determine how to bid across auctions. Even where the bids are collected at the same time, bidders still need to guess how to bid.⁹

 $^{^9 \}mathrm{See}$ [56]. It is unclear from Mueller whether the auctions were simultaneous or sequential.

As noted above, there are incentive-compatible VCG designs for multi-product auctions. Chapter 10 below characterizes prices and outcomes in such auctions, as well as the variants of VCG auctions being used for selling spectrum licenses in a number of European countries.

4.4.4 VCG-like Mechanisms for Advertising Auctions

One area in which a variant of the second-price auction is commonly used is for position in the ad sections of online search engine web pages, such as Bing and Google. When someone enters a keyword, e.g., "plumber," the top of the screen will contain a small number of ads. Which ads appear depends on continuous online bidding for the positions. These auctions now generate billions of dollars in annual revenues.¹⁰ The following describes the mechanism used for these auctions, which is a variant of the VCG auction.

To use Varian's [71] notation, it is assumed that there are A agents competing for S < A slots. The bids typically take the form of a fee paid by the winner per *click-through*. The higher slots are assumed to have a higher *click-through rate* (ctr), x_s . So, a bidder that derives a profit of v_a per click will earn a profit of $u_{as} = v_a x_s$ from winning slot s. The slots are numbered in such a way that $x_j > x_{j+1}$ for all j = 1, 2, ..., S-1. Table 4.5 summarizes the Google auction.

Position	Value	Bid	Price	CTR
1	v_1	b_1	$p_1 = b_2$	x_1
2	v_2	b_2	$p_2 = b_3$	x_2
3	v_3	b_3	$p_2 = b_3$	x_3
4	v_4	b_4	$p_1 = b_4$	x_4
5	v_5	b_5	$p_1 = 0$	0

Table 4.5: Position Auction Bids and Payoffs

Bids for these slots take the form of a bid per click-through. So, if a bidder with value per click-through v_a pays $p_s = b_{s+1}$ for slot s, then that bidder's profits are $x_s(v_a - p_s) = x_s(v_a - b_{s+1})$. In other words, the bidder that wins

¹⁰Varian [71] gave a figure of over \$10 billion for 2005, just for Google and Yahoo.
4.5. SUMMARY

slot s pays the value the next highest bidder would have derived from that slot.

Varian shows that this pricing rule will result in the same allocation and payments at a symmetric Nash equilibrium as the VCG pivot mechanism.¹¹ Thus, the position auction does provide bidders some incentives to bid truthfully. However, this result rests on a number of assumptions. I briefly list a few.

First, it is assumed that bidders each would receive the same ctr, x_s , from a given slot. It is unclear why this assumption is reasonable. Indeed, bidders may place different priorities on getting the top slot.

Second, bidders may have advertising budgets limiting what they can bid. Suppose, for example, there are two bidders, each having a budget of 100, and each wanting to buy a high position for the same three keywords. There will be no way for two bidders to allocate their budgets so that the resulting outcome is a Nash equilibrium in pure strategies. Each will want to try to pick up one slot cheaply, leaving a lot over for a second slot. Indeed, there is no equilibrium in pure strategies.

Third, values of some slots can depend on other slots won. There can be both substitutes and complements. The Google auction does not allow for the type of contingent or package bids that are required if bidders are to convey true values.

4.5 Summary

This section concludes with a brief discussion and example of how the principles of optimal auction design discussed in this chapter can be put into practice. As is explained below, there may not be a single optimal auction design, or if there is one, the conditions under which it applies may be overly restrictive. This discussion uses an example from an actual auction design problem to illustrate some of the main issues.¹²

¹¹The pivot mechanism has each bidder pay the difference between the aggregate value of rivals' winning bids without that bidder and the value with that bidder.

¹²Much of what follows is based on [53] and joint work done on a related project.

The design problem being addressed was to determine how to best allocate universal funds to support universal telecommunications service goals in the United States. The US Federal Communications Commission (FCC) and the state regulatory commissions established the Joint Board on Universal Service. The purpose of this agency is "to make recommendations to implement the universal service provisions of the Act. This Joint Board is comprised of FCC Commissioners, State Utility Commissioners, and a consumer advocate representative."¹³ The Act established a fund to promote deployment of telecommunications services in rural areas in for under-served segments.

The problem was to find an efficient auction design. This design had to meet the objective of minimizing costs for each market area receiving funding. In other words, the auction design was to maximize the benefits that could be provided with a fixed available amount of universal service funds. The auction design needed to take account of a number of features of the market. First, there was the problem of dividing the market into areas of a size appropriate to the potential service providers. This is a problem addressed in more detail below. For universal service funds, the potential service providers included wireline telephone companies, cable television, and cellular firms. This suggested certain geographic areas.

One additional concern was that there could be post-auction holdup or performance failures. Auction winners would be more likely to meet their universal service commitments if there was competition in the market. Unfortunately, this form of competition can result in excessive costs due to unnecessary duplication of facilities. The proposed way to address this tradeoff was to include in the benefits accruing to consumers an extra factor for the presence of post-auction competition in a market. This meant that a second provider would be selected when the costs imposed by that provider were below a certain level.

The auction design problem was to derive conditions on the set of winners and payments that had to be satisfied at an optimum. What was shown was that an optimal auction design was any design for which the efficient set of providers almost always won, and for which the highest-cost winners earned zero profits.¹⁴ What this result does not provide is an explicit char-

 $^{^{13} \}rm See~http://transition.fcc.gov/wcb/tapd/universal_service/JointBoard/welcome.html. <math display="inline">^{14} \rm See~[53]$ for details.

acterization of the optimal auction design; rather, it indicates than there is an entire class of auction designs that all achieve the same outcome. This general result does have some directly useful implications. For instance, it indicates that the market structure should be endogenous. However, it is left for the auction designer to develop specific rules and verify they do indeed satisfy the conditions for an optimum allocation.

Chapter 5

Imperfect Information and the Winner's Curse

Summary

This chapter describes auctions in which bidders have imperfect information about the value of the object for which they are bidding. One specific case is that in which the object has the same common value for each bidder, but bidders have different ex ante estimates of the value. This chapter also examines auctions in which bidders have affiliated values, that is, bidders' values are correlated in some way. What this means is that winning provides information, and so bids should be adjusted to allow for this fact.

5.1 Common Values and Order Statistics

Competing bidders often have to bid without having perfect certainty over values. A bidder may have an ex ante belief, or prior, about the value. That prior, which can be modeled as a probability distribution, and one in which the possible range of values is bounded both above and below, may be correlated across bidders. For example, in any auction for a durable asset with a resale market, such as real estate or a building, bidders' values will be correlated, as part of each bidder's value is the potential gain from resale. In the *pure common-value* case, all bidders will have the same value ex post, even if they have different forecasts of value ex ante. Bidders for oil leases will usually have only imprecise, and sometimes asymmetric, information about the amount of oil available in the tract being auctioned,¹ but all bidders may have the exact same value ex post – the value of the oil in that tract. In any auction in which value depends on future prices, such as oil for oil leases, or on complementary products, the bidders' valuations must contain some common-value components.

In a common-value auction, the winning bidder will tend to have the most optimistic forecast value, and thus tend to overestimate the value. This will be true, for instance, when each bidder has access to a signal of the value of an object; then the bidder with the highest signal will generally want to bid the most. In this case, the winning bidder will tend to overpay whenever the underlying value is the same for all bidders, i.e., in a pure common-value auction, unless that bidder adjusts its bids to allow for the fact that being high bidder means it has the most optimistic signal. If each bidder bids up to its ex ante, unconditional expected value, as it would in a second-price auction in the pure private values case, the winning bidder will likely bid too much.

To see why this is the case, suppose that the true value of an object is V, and n bidders have signals s_j , j = 1, 2, ..., n, uniformly distributed on $[V - \frac{k}{2}, V + \frac{k}{2}]$. If bidder j offers s_j , then the expected price is $\frac{V+kn-3}{2(n+1)}$, which is greater than V whenever there are four or more bidders. In other words, the bidder with the highest signal will have a value which, on average, will exceed the true value by $k\frac{n-1}{2(n+1)}$. Knowing this, such a bidder should discount its bid accordingly. A naive bidder, failing to make such an adjustment, might win the auction, but suffer losses. Such was the experience of bidders for oil leases as documented in [17].

It is easy to see that the RET can fail when there are common values, or even when bidder valuations are affiliated, that is, imperfectly correlated. Consider the above example of a common-value auction. In a commonvalue auction, the rival bidders can have accurate forecasts on average, but the winner in each auction will tend to be that bidder that has the most

¹See [17] or [35] for a discussion of oil lease auctions.

optimistic forecast in that auction. In the second-price auction, bidding the ex ante signal of its valuation is the dominant strategy when values are independent. However, while it still might be the case that a bidder should bid its ex ante value in a second-price auction, conditional on being the winner among n bidders, it is no longer true that this value should be the stopping price in the English auction is adjusted each time a different bidder drops out. The question addressed here is how auction formats compare with common or affiliated values.

This chapter also addresses the more general question of how imperfect and asymmetric information can affect bidding behavior and the auction outcome. More specifically, the precise information structure can affect bidder behavior. When bidders have affiliated values, the structure of information can matter a great deal. As is explained in further detail below, it can matter more whether information is private than whether it is precise.

5.2 Bidding Strategy in Common-Value Auctions

This section assumes a pure common-value auction, that is, each bidder *i* has a signal $x_i = v + \epsilon_i$ of the true value v.² It is also assumed that the signals are drawn from a common density function $f(\epsilon)$. Note that the expected value, conditional on all *n* bidders' signals, is $\frac{1}{n} \sum x_j$. While it is true that $E(x_j) = v = \frac{1}{n} \sum x_j$ (where E(Z) denotes the expected value of Z), it is not the case that $E(X^1) \equiv E(\max_j \{x_j\}) = v$. In other words, the likely winning bidder will usually be the bidder that has the highest signal and will have an expected value, conditional on winning, that does not equal the unconditioned expected value. Thus, in an auction, the bid strategy must be adjusted for this bias, to avoid the *winner's curse*. This adjustment will be based on the information available to the winning bidder, and that bidder has to make his or her final decision about a last and best offer. Because the information available can depend on the type of auction, the expected

 $^{^2 {\}rm In}$ contrast, in a pure private values case, each bidder's valuation is independent, i.e., there is no common value term, v.

revenues of different auctions, such as English and Dutch, need not be the same.

5.2.1 Bid Strategy in Dutch and First-Price Auctions

In a first-price auction, (4.1) determines the optimal bidding strategy, assuming independent private values. As above, let $\beta(s)$ denote the equilibrium bid function for all other bidders, and assume a symmetric equilibrium. In the case of common values, it is still true that a bidder with a signal s of value v will, at equilibrium, not want to shade its bid below (increase its bid above) $\beta(s_i)$ when the lower (the higher) probability of winning will not exceed the benefits of reducing (increasing) its cost should it win. That is, using the notation of Chapter 4,

$$E\left[v - \beta(s_i + \Delta s)|s_i + \Delta s \ge s_j, j \ne i\right] \le E\left[v - \beta(s_i)|s_i + \Delta s \ge s_j, j \ne i\right]$$

This condition just states that a bidder with signal s_i will not want to report it has a different signal $s_j \pm \Delta$. Moreover, this condition is the same as used in deriving (4.1) except that the expectation was not conditional on the bidder having the highest value signal. More to the point, in the independent-private-values case of Chapter 4, the probability that a bidder would attach to winning is not affected by others' values, whereas bidders' values are so affected in the common-value model. Thus, it is still the case that the optimal bidding rule has the same form as in (4.1), with the exception that the distribution function is conditional on the bidder in question having the highest value, i.e.,

$$\beta(v) = v - \frac{K + \left\{ \int_0^v F^{N-1}(s|v>s)ds \right\}}{F^{N-1}(v|v>s_j \text{ for all other bidders } j)}$$
(5.1)

Moreover, it is also the case that a bidder in a first-price auction will face exactly the same tradeoffs as in a Dutch auction. In both cases, each bidder will want to shade its bid a bit, that is, bid below value. Moreover, in each case a bidder will face the same tradeoffs in making that decision. A winning bidder will want to enter a bid, at an amount less than value, that maximizes its expected surplus conditional on its having the most optimistic value. Thus, the strategic equivalence of the Dutch and first-price sealed-bid auctions carries over to the common-value case.

5.2.2 Bid Strategy in a Second-Price Auction

In a pure common-value auction, the firm with the highest value estimate will win and pay the amount offered by the bidder with the second highest value. So, if any bidder wins, it will have the highest signal, and its offer should be discounted accordingly. The offer should equal how much the bidder expects the value to be, conditional on its forecast having the highest value. The amount that a bidder should offer is still equal to its true *expected* value. However, the expected value is no longer the signal or forecast value.

5.2.3 Bid Strategy in English Auctions

When bidders' valuations are independent, the optimal bid strategies in the second-price and English auctions are essentially equivalent. In the English auction, a bidder will want to bid up to its forecast value, and then drop out; in the second-price auction, a bidder will bid its ex ante forecast value. The outcome is also the same. To see this is no longer true when there are common values, it suffices to consider the decision facing a bidder about when to drop out in an English auction.

Consider a bidder that has a signal s about the value of the object being auctioned. Further assume that this is a pure common-value auction. What follows describes the bidding decision for a Japanese-style English auction, that is, an auction which starts with a low price which the auctioneer gradually increases until there is only one bidder remaining. Bidders, at each increment, must indicate that they are still in the auction or drop out. Once a bidder drops out, it is never allowed back. And bidders each see how many bidders remain after each bid increment.

If the bidder has not seen any rival drop out, then the worst case – that is, the case in which the bidder would have the lowest conditional forecast value – is when all rivals have a signal which is just the same as its own. So, if the true mean forecast is $\frac{1}{n} \sum_{k=1}^{n} s_k$, a firm with the signal s should drop when the price reaches the expected value $E[V|s_k = s \text{ for all } k]$.

Let $p^1(s)$ denote the value of the first drop. Notice that this is a strictly lower amount than a bidder might forecast to be the expected value, assuming it has the lowest value signal.³

Then the next drop must occur when the bidder with the second lowest value signal can expect to lose money if it should stay in. The decision about when to drop for the bidder with the second lowest value signal needs to be based upon what its forecast would be, conditional on its winning. The only way such a bidder (call it bidder 2) can win is if all its remaining rivals have the same value as it does, unless it bids too much. Bidder 2 will bid too much if it stops after higher-value bidders have dropped out. The worst case for bidder 2 occurs when every rival has the same value, just a bit more than bidder 2's value, and they all drop out at a price just above where bidder 2 might. Bidder 2 can wait, with the expectation that the actual value is better than the worst case; however, this bidder does not want to win if it actually has the second lowest signal, and so its choosing to drop out, assuming the worst case, is, in this instance, a dominant strategy.⁴ Let $\beta_k(s, h_{k-1})$ denote the strategy for the kth bidder to drop out given that the history of the k-1previous drops is h_k and the bidder has signal s. Also, suppose bidders are indexed in order of the value of their signals.

Let $\beta_2(s, p^1) = E[V|\beta_1(s) = p^1(s), s^j = s \text{ for all } j \ge 2]$, where s^j is the *j*th highest signal. Then $\beta_2(s, p^1)$ specifies the optimal price at which to drop out of an English auction for a bidder with the second lowest signal. One can iteratively define an optimal stopping price for each bidder,

$$\beta_k(s, h_{k-1}) = E[V|s^j = s \text{ for all } j \ge k, h_k]$$

These strategies yield optimal prices for a bidder to drop out, except for the last two bidders. The last two bidders will see that all other bidders have dropped out. Thus, the remaining bidder with the lower signal should

³Suppose, for example, forecasts are uniformly distributed on some interval $[V - \frac{1}{2}, V + \frac{1}{2}]$, where V is the unknown true value. Then the expected value of V for a bidder that knows it has the lowest forecast value s^N among N bidders is $s^N - \frac{1}{N+1} + \frac{1}{2}$.

⁴If there are only two bidders, then this analysis needs to be modified, as in this case the second lowest bidder would want to wait longer.

assume its last remaining rival has a higher signal. For example, suppose that there are two bidders to start the auction. Also, suppose that these two bidders have signals that are $s^j = V + \epsilon_j$, where ϵ_j is uniformly distributed on $\left[-\frac{1}{2}, +\frac{1}{2}\right]$. If the bidders do not know the true value V, then the lower-value bidder should drop out when the price reaches its signal plus $\frac{1}{6}$, as the expected value of the lower value signal will be $V - \frac{1}{6}$, and that of the higher value signal will be $V + \frac{1}{6}$.

This induction approach defines optimal bid strategies for a sequence of common-value English auctions. Assuming risk-neutral bidders, bidders will drop out in the order of the value of their signals, and the bidder with the highest signal will win, but will only pay the price at which the secondhighest-value bidder drops out.

5.2.4 Comparisons between English, Second-Price, and Dutch Common-Value Auction Outcomes

The above provides an analysis of the optimal bidding rules in English, Dutch or first-price, and second-price auctions when there are common values. Common values mean that bidders need to adjust their bids to reflect that winning provides information. In other words, the winner will tend to have the most, and an overly, optimistic forecast.⁵ In many auctions, bidders' values have common-value components. [17] observed that bidders for oil leases often failed to adjust for the winner's curse, and as a result winning bidders tended to overbid. These were not pure common-value auctions, i.e., bidder values included not only the common value of the underlying oil, but also the knowledge, costs, and ability to exploit the oil lease, which could differ across bidders. However, bidders still needed to adjust for the likelihood that winning means that a bidder has made an overly optimistic forecast.

As an illustrative example, suppose bidder signals of values are distributed uniformly on some interval $[V - \frac{1}{2}, V + \frac{1}{2}]$, and bidders do not know V. Suppose that the true value is V. So, the expected value of the *j*th lowest signal will be $s^j = V - \frac{1}{2} + \frac{j}{n+1}$.

⁵If there are only two bidders, or if the number of lots available is more than half the number of bidders, then the marginal losing bidder will tend to have a forecast value that is below the true value.

In a second-price auction, a bidder that has a signal s will have expectation that the true value is $s - \frac{n-1}{2(n+1)}$ conditional on its being the winner. This implies that the expected value for the bid amount of the bidder with the second highest signal is $V - \frac{n-1}{n(n+1)}$.

On the other hand, in an English auction, the bidder with the *j*th lowest signal will want to bid up to the point at which the price equals $\frac{s^1+s^j}{2}$. The expected value of this price is then $\frac{V}{2} - \frac{V}{2(n+1)}$. That is, in the English auction, the expected price is higher, and closer to the expected value. This is not always true, but when there are common values, the expected price in the English auction cannot be less than in the second-price sealed-bid auction.

And finally, in a first-price sealed-bid auction, a bidder with signal s that offers $\beta(s) = s - x$ will have a probability of winning of $s - V + \frac{1}{2}$, and will win V - (s - x). The bidder will want to maximize the expected value of this with respect to x. It can be shown⁶ that the optimal bid is $\beta(s) = s - \frac{1}{2}$, or $x = \frac{1}{2}$. This results in an expected value for the price of $V - \frac{1}{n+1}$.

Summarizing, the expected price is

- In an English auction: $V \frac{1}{2(n+1)}$
- In a second-price auction: $V \frac{n-1}{n(n+1)}$
- In a first-price auction: $V \frac{1}{n+1}$.

So, in this example, and more generally, the equilibrium price in an English auction will be the price from a second-price auction, which, in turn, will exceed the price from a first-price auction.

5.3 Almost Common-Value Auctions and Informational Asymmetries

Bidders can differ in the amounts of information they have access to and in their underlying values. For example, in the types of oil lease auctions

 $^{^{6}}$ See [41].

described in [17], some bidders may have access to better geological information than their rivals. Moreover, there are investments in complementary facilities that some bidders may have already incurred and rivals have not.⁷ The optimal auction strategy, and thus the outcome, will depend a great deal on how information is distributed among the bidders.

The preceding section considered the case in which, ex ante, all bidders have the same quality of information, although not usually identical information. Bidders may all derive value from the same unknown variable, such as the amount of oil in an oil tract, but will typically have different views or forecasts as to the magnitude of this common value. This type of model is sometimes referred to as an oil lease model.

However, bidders need not be symmetric. First, some bidders may be better informed than their rivals. This is the situation in drainage or resuscitation tract auctions; in such auctions, those bidders that have previously explored or developed the tract will have knowledge about reserves and costs that rivals do not have. Sometimes, informed bidders may not all have access to the same information. In other cases, bidders will have some private values, that is, economic advantages. This is the case, for instance, when some bidders have previously invested in complementary assets, such as storage and pipeline facilities. This cost advantage, or value premium, gives these bidders an advantage over rivals.

5.3.1 Auctions with One Informed Bidder

This subsection considers the situation in which one bidder is informed, and the other bidder or bidders are uninformed, about the value of the object for sale. In the simplest case, the value can be high, H, with ex ante probability ρ , and low, L, with probability $1 - \rho$. Assume too that this is a first-price sealed-bid auction. In this situation, the informed bidder will want to bid just high enough to make uninformed bidders indifferent between bidding and not bidding. More specifically, let $x = \rho z H + (1 - \rho L)$, where z is the probability that an uninformed bidders wins when the value is actually H. If the uninformed bidder always bids x, and the informed bidder offers x

⁷See for example [35].

when the value is H, then the uninformed bidder will be indifferent between bidding and not. Any lower bid will mean that the uninformed bidder can bid a bit more and always win when the value is high. Any higher bid will mean the uninformed bidder would not bid.

This logic can be generalized for any number of uninformed bidders. Somewhat surprisingly, the outcome is the same. One key requirement for an equilibrium is that uninformed bidders must just break even, conditional on winning; otherwise the informed bidder should raise its bid. If there is only one informed bidder, then the equilibrium bid is the same no matter how many uninformed bidders there are. The other key requirement for an equilibrium is that no bidder has an ex ante negative expected payoff.

Notice that a pure common-value English auction ends up as essentially a first-price sealed-bid auction for the two bidders that have the most optimistic forecasts. After all but the last two bidders have dropped out, each of the two remaining bidders will know its own signal, and each will know that, if it is not the last to drop, then it will have the highest value signal. The one with the high value signal will know this value, which its only remaining rival can only guess at. Analytically, the firm with the lower forecast of the two will not want to wait too long to drop. The difference is that in the common-value auction, the last two bidders do not know which is the stronger until after the auction ends.

5.3.2 Auctions with Both Informed and Uninformed Bidders

When there are two or more informed bidders, and uninformed rivals, the uninformed rivals will generally be at a critical disadvantage. If the informed, or better-informed, bidders have identical information, then they will bid prices down to the point where uninformed bidders cannot break even. In other words, uninformed bidders do not want to win. For any common signal, the informed bidders will never, in equilibrium, submit any offer that will leave a positive profit. If one informed bidder did submit such an offer, then an informed rival would seek to submit a slightly lower bid amount. Thus, if an uninformed bidder wins, it will have bid too much, and will lose money. This situation is then logically equivalent to a pure common-value auction among the informed bidders, with uninformed bidders not bidding.

The situation is different when bidders have different information. In that case, if several bidders have access to the same information, then none of those bidders will be able to earn positive expected profits. On the other hand, if one bidder has access to even very imperfect signals about the object's value that its rivals do not have, it can earn a profit. For example, suppose one bidder has a signal of value that has a high variance, and several other bidders have access to an identical, low-variance signal. The bidders that have the same information will compete away profits. When the bidder with different information has a positive signal, while the others have a negative or lower signal, it can win, and earn a positive return. In other cases it will not win.

5.3.3 Other Asymmetries

Most auctions have common-value aspects. Often bidders differ in many other ways. After any uncertainty is resolved, bidders can have much different underlying values – due, for instance, to differences in costs or ability to derive benefits from what is available. If a single bidder has a value or cost advantage as well as an informational advantage, then the analysis is essentially the same as in the case in which there is a single informed bidder competing with one or more uninformed bidders.

On the other hand, with several informed bidders, only one of which has a value or cost advantage over its rivals, the higher-cost informed bidders should not win, but their presence will mean that uninformed bidders should not bid, or bid only very low. An uninformed bidder can have a cost advantage over an informed rival. This case is more complex in that the information advantage can fully or partially offset the benefits of the value differential.

More specifically, the informed bidder may choose to bid low, and accept a low or zero probability of winning, if it has a negative signal of the value, and to bid aggressively if it has a positive signal. The bidder with the higher value may want to reduce its bid, knowing it will derive lower net value if it wins. If it may be the case that the strong bidder will never want to allow the informed bidder to win: it will want to keep its bid high enough to win all the time. Alternatively, the strong bidder may want to allow the informed bidder to win when it has a positive signal. In this case, the informed bidder will bid just high enough to top the informed bidder when the value is low.⁸

5.4 Summary

This chapter has examined the influence that uncertainty in a common or affiliated value has on the outcome of an auction. Bidders facing this type of uncertainty will want to adjust bids to offset the impact of the winner's curse: winning means a bidder has the most optimistic forecast, which will usually be too high. One of the key results of this chapter is that an open auction allows bidders to aggregate information. This reduces the winner's curse and increases auction revenues. Thus, an English auction will generate greater revenues than a first-price or second-price sealed-bid auction.

This chapter also explains how one can calculate an optimal bid strategy to correct for the winner's curse. The calculation essentially involves correcting estimates of values and optimal bids for the fact that having the winning bid means that the bidder has the highest (most optimistic) signal. It turns out that the additional information available to the winning bidder in an English auction increases the precision of value forecasts. The second-price auction will also generate more revenue than a first-price auction.

This type of analysis has had some practical and observable implications for oil lease bids. This is one situation in which bidders have experienced the winner's curse, and presumably learned how to adjust their bids. The empirical evidence⁹ suggests that bidder experience has allowed bidders to adjust bids in line with what is an ex ante optimal bid strategy. [30] also suggest that this is also likely to be the case with asymmetric information.

⁸Suppose the strong bidder places a value that is 2 higher than the uninformed bidder's, and that the value to the informed bidder can be 2 or 8. If the probability of a high value is sufficiently great, viz., 2p - 4(1-p) > 2(1-p), where p is the probability of a high value, then the strong bidder will always win, and will bid 8. This bidder will win 2 with probability p and lose 4 with probability 1-p. For smaller values of p, the strong bidder will prefer to allow the informed bidder to win when the value is high.

 $^{^{9}}$ See [60].

5.4. SUMMARY

Undoubtedly, the winner's curse arises in other auctions, but it is less well documented than it is for oil lease bids. There is, however, extensive experience with spectrum auctions, and there are numerous cases in which a bidder has won only to find it overpaid. The most striking example is the German 3G auction in 2000. In that auction, two bidders, Group 3G (a consortium which included Sonera and Telefonica) and Mobilcom Multimedia (with France Telecom as the major mobile operator in the consortium) each paid approximately \$8 billion for a license, only to abandon it a short while after the auction. Neither operator ever launched.¹⁰ These were very large, and presumably well-prepared and sophisticated, firms.

The ability of bidders to offset for the winner's curse can be very limited. Experimental evidence suggests that the calculations can be hard for bidders to make accurately. As the theory suggests,¹¹ providing more information to bidders should improve welfare and efficiency. In practice, however, it is unclear in many cases (especially one-off auctions or auctions that are very sporadic for bidders) how to correct accurately for the winner' curse.

 $^{^{10}}$ A few years later, Telefonica reentered the German market through an acquisition of one of the *winners*, and was losing money as of 2011.

90CHAPTER 5. IMPERFECT INFORMATION AND THE WINNER'S CURSE

Chapter 6

Sequential Auctions of Substitutes

Summary

This chapter examines sequential auctions of multiple lots of an identical good or service. The question addressed is whether price systematically increases or decreases from one auction to the next. Such systematic price patterns can be an important consideration for bidding strategy and for auction design. First considered is a sequence of auctions with an identical pool of risk-neutral bidders. It is shown that under fairly general conditions the expected price will tend to neither increase or decrease across auctions. It is also shown that when bidder opportunity costs can change from one auction to the next, or when bidders are risk-averse, definite pricing trends will arise.

6.1 Multiple-Lot Auctions

A single auction or auction event will often include multiple units of identical or similar products. This chapter examines auctions in which the lots are sold sequentially, generally one at a time. Subsequent chapters consider alternative auction designs in which a single auction process, or even a single package bid, can include multiple units, or even multiple objects. 92

Bidders participating in a sequence of auctions will need to decide how high to bid in any given auction when there are other auctions to follow for identical or similar objects. This is especially the case In a sequence of auctions; an aggressive bid early in the sequence can lead to regret when prices are significantly lower in subsequent auctions. Conversely, a bidder who decides not to bid up to his or her value in an early auction can be disappointed with much worse prices, and the possibility of being shut out, in subsequent auctions.

This chapter describes solutions to the sequential decision problems facing the originator and the bidders. The analysis of the solutions to these decision problems is then applied to characterize expected equilibrium outcomes in different types of sequential auctions. Most of what follows assumes that each bidder is interested in at most one unit, so that strategic withholding is not a consideration. Thus, a bidder considering whether to improve its offer for any lot will have to weigh the extra probability of winning the current auction against expected surplus that it will receive for waiting for one of the next lots. The number of remaining lots and the number of bidders can, in theory, affect this calculation, and in a surprising way. This calculation can be affected if bidders are risk averse, or if their opportunity costs change across auctions.

The auctioneer, at times, will also have a choice about how to allocate what is available for auction over time and across auctions. This type of timing decision is common in energy procurement¹ and in many other sectors.² As is explained below, this decision can affect overall revenues. In addition, the auctioneer may consider different reservation prices, and different ways to split the quantity available across auctions. The auctioneer can consider reservation prices based on offers in the auction, choosing to defer some volume to the next if competition is low, and the reverse if competition is high, in an initial auction; in other words, the auctioneer can choose the quantity to carry over from one auction to the next, based on the level of competition as well as on other factors, such as pre-auction price expectations.

¹See for example www.bgs-auction.com. These auctions stagger purchases over time. In addition, each auction contains a provision for deferring procurement to a later date should participation in any one auction be limited. See also [49].

²See for example www.globaldairytrade.info for wholesale dairy auctions, and www.cranberryauction.info for wholesale cranberry auctions.

These decisions become more complex when the products are imperfect substitutes. In this case, the order of what is offered can affect bidder decisions and the overall outcome.

6.2 The Declining-Price Anomaly

An item sold in one auction will often have a copy, or a similar item, available in subsequent auctions. It has been observed that prices will tend to vary quite a bit even when identical objects are auctioned one after another in the same auction event. This suggests that auction design and strategy can affect the outcome. This also appears to contradict one of the fundamental results of economic theory – that all parties face the same price for the identical product in a competitive market. [4] refers to this lack of price uniformity in these auctions as "the repeal of the law of one price." This phenomenon has been the subject of numerous theoretical and empirical studies.

This lack of consistency in prices was noted in a study of wine auctions and is illustrated in Figure 6.1.³ The first auctions for the 1961 Palmer and 1927 Croft each sold for a significant premium with respect to the final auctions. This data appears inconsistent with standard models of perfect competition. This and other empirical evidence suggests what is often called an *afternoon effect* or a *declining-price anomaly*; that is, as in these wine auctions, for the early lots to sell at a premium relative to the later lots is fairly common, and much more common than for prices to increase across auctions.⁴ The task addressed in what follows is to identify the factors that determine what patterns of prices arise in practice.

In theory there are two effects. One is that bidders should want to bid less aggressively for the first lots than for the later lots. In the first auction there are more remaining chances to win. In the last auction, bidders should be willing to bid up to their reservation values, but not in the other auctions. This logic means that a bidder's optimal bidding strategy should be to bid an increasing fraction of its reservation price from one auction to the next. However, the countervailing factor is that the first winners will tend to be

 $^{{}^{3}}See [4]$.

 $^{{}^{4}}See [4] \text{ and } ([5] .$

	Christie's London	Sotheby's London	Christie's Chigaco	Butterfield's San Franc.	Total
Later price higher	271	143	90	20	524 (11%)
Later price lower	628	430	183	41	1282 (28%)
Later price	1498	1073	226	39	2836

Distribution of price patterns for identical wines sold in same auctions (number of auctions)

	Chateau Palmer 1961				Croft 1927		Cl Marg	Chateau Margaux 1952		Quinta de Noral 1934		
	Lot size	Price	Price/ bottle	Lot size	Price	Price/ bottle	Lot size	Price	Price/ bottle	Lot size	Price	Price/ bottle
Lot 1	12	920	77	12	800	67	12	480	40	10	400	40
Lot 2	12	800	67	12	800	67	12	480	40	12	500	42
Lot 3	12	700	58	12	750	63	12	480	40	12	500	42
Lot 4				12	650	54	24	480	20	12	480	40
Lot 5				12	650	54	24	480	20	12	480	40
Lot 6				12	650	54	20	480	24			
Lot 7				12	650	54						

Source: "Liquid Assets", The International Guidet o Fine Wines, I ssue No. 4, Spring 1988.

Figure 6.1: Sequential Wine Auctions

those with the highest valuations, so competition should tend to decrease from one auction to the next. In theory it is possible for these two effects to cancel out. One main result of this chapter is that under certain conditions, they do. This chapter also explains some conditions under which price can be expected to increase or decrease.

This chapter also explains when prices can be expected to increase or decrease across auctions. Briefly, when a seller (a buyer) holds two equivalent auctions, then the buyers (sellers) bidding in the first auction will want to offer a slightly lower (higher) price than the expected value of the price in the second auction. The intuition behind this result is that a bidder in the first auction will want to pay no more than the expected price in the second auction less an amount to adjust for the expected surplus from winning the second.

6.3 Weber's Martingale Theorem

This section presents a fundamental and surprising result, Weber's *martin-gale theorem* characterizing the expected values of prices in a sequence of auctions. This result states that the expected price in one auction is, under conditions explained below, equal to the price in the previous auction. In other words, the most recent auction price is the best predictor of the price in the next auction. The intuition behind this result derives from the fact that if strategic bidders were to think one auction would produce a lower price, these bidders would all want to participate, and remain active, in that auction, until the price reached the level of other auctions. In this case, the auction design, and more specifically the allocation of the total quantity across auctions, should not affect the expected price in any auction or the overall auction revenues.

Consider a model in which there are some number B of bidders, and I < B items for sale. Each bidder can win at most one lot. To simplify the exposition, it is assumed that each bidder has a valuation v_b and that the bidders' valuations are uniformly drawn from a uniform distribution over the unit interval. Assume bidders are labeled so that $v^b > v^{b+1}$ for all $b = 1, 2, \ldots, B - 1$. These assumptions allow an easy characterization of

equilibrium, and as will be clear from what follows, the results apply quite generally. What follows assumes that the lots are sold in a sequence of firstprice auctions.

The following basic insights allow a direct calculation of the solution to the auction.

- 1. Backward Induction: To solve for the equilibrium of this auction it is easiest to start working backwards, from the last auction. First, assume all but one item has been sold, and that there are B I + 1 bidders remaining.
- 2. Revenue Equivalence Theorem (RET): A consequence of the RET is that the highest-value remaining bidder, I, should win the last auction and have to pay a price equal to the value of the second highest remaining bidder. So, the price that the high-value remaining bidder pays should equal v^{I+1} , the value of the I + 1st remaining bidder.
- 3. Order Statistics: The expected value of v^b , the *b*th highest valuation or the *b*th order statistic, is $\frac{B+1-b}{B+1}$. This means that the expected value of v^{I+1} is $\frac{B-I-1}{B+1}$. Notice that this ratio is close to 1 if *I* and *B* are large and *I* is almost as large as *B*. If $I = \alpha B$, $\alpha \in (0, 1)$, and both *I* and *B* are large, then the price will be approximately $1 - \alpha$.

The above implies that last item sold will sell for a price of $\frac{B-I-1}{B+1}$, assuming this last auction is a second-price sealed-bid one. (Note that the RET implies that the item would also sell for this same price if the last auction were a first-price sealed-bid or an oral ascending auction.) Bidder I's value will on average be $\frac{B-I}{B+1}$, and I will pay $\frac{B-I-1}{B+1}$, which is the fraction $\frac{B-I-1}{B-1}$ of I's value.

Then, the question is what should be the price of the next to last item sold. Assuming that the I-2 highest-value bidders win items $1, 2, \ldots, I-2$, then the remaining bidders will be $I-1, I, I+1, \ldots, B$. Bidder I should win that last auction, but not pay more than $\frac{B-I-1}{B+1}$, the same price as bidder I. Therefore, in the next last to auction, bidder I-1 should win on offering a price which equals the fraction $\frac{B-I-1}{B-I+1}$ of its value.

6.4. STRATEGIC ALLOCATION DECISIONS

By backward induction, all I objects sell for the same price $\frac{B-I-1}{B+1}$, and the bidders should offer increasing fractions of their value, $\frac{B-I-1}{B-1}, \ldots, \frac{B-I-1}{B-I+1}, \frac{B-I-1}{B-I}$.

In this example, the expected price is constant, and in particular, the expected value of the price in each auction is the same as the price in the last auction. This property means that the price is what is called a *martingale*. This result is quite a bit more general.

Proposition 1 Suppose there are I identical items available in a sequence of first-price auctions. Suppose that each bidder can purchase one item, and that a bidder with valuation v will bid more for an item in an auction than a bidder with valuation w < v.

Then, the expected value of the price in each auction will equal the price in the previous auction, i.e., $E(p_i|p_{i-1}, p_{i-2}, \ldots, p_1) = p_{i-1}$.

Sketch of Proof

Let I_j denote the information available when item j + 1 is to be sold. Notice that RET implies that the expected payment by winners of the last I - jlots is $E[v^{I+1}|I_{j+1}]$. Then the expected price as of auction j for items sold in subsequent auctions $j + 1, j + 2, \ldots, I$ is $E[E[v^{I+1}]|I_{j+1}|I_j] = E[v^{I+1}|I_j]$. Therefore, this price must be the same as the price for auction j. \Box

Notice that the martingale theorem extends to the case in which the auctioneer can divide the available volume across any fixed number of auctions. More specifically, suppose that the auctioneer conducts a sequence of uniform-price auctions. In each auction a, l_a lots are sold. Also, suppose that in each auction the top l_a bidders win and pay the lowest of the winning offers. Then, a similar argument implies that all items will sell for a price equal to v^{I+1} , independently of the how the items are distributed across lots.

6.4 Strategic Allocation Decisions

In practice, auctions can include multiple units sold (or purchased) at one time. Thus, an auction originator can face the decision how to divide the amount available across auctions. The martingale theorem suggests that decisions about how to allocate quantity sold (or purchased in a reverse auction) across auctions should not affect the expected outcome.

There are a variety of reasons why this is not necessarily the case. One of the most common is uncertainty about overall market conditions. For example, in a series of ongoing energy procurement auctions, the utilities conducting the auction leave open the option of limiting procurement volume if transitory market conditions are highly unfavorable. Transitory energy price spikes that occur on the day of an auction could leave the utilities, and the ratepayers who are their customers, paying prices that reflect the price spikes, not the underlying market conditions. Unusual events, such as hurricanes, can create significant short-term price uncertainty, which can justify some auction volume reductions.⁵

This section explains two other factors: one is how the split of the auction volume across time can affect relative competition, bid strategy, and prices across auctions, and the other is how resolution of uncertainty may affect the outcome.

These questions are related to the question of the strategy a buyer or seller should take when choosing between forward and spot market transactions. [3] consider a model in which firms can compete in both spot and forward markets. The offers in the forward market affect the payoffs in the spot market. This can be modeled as a two-stage auction – the first-stage sales are not made to the highest-value bidders, but to a subset of *average* buyers over the two stages. What Allaz and Vila find is that if one firm participates in the forward market and the other cannot, then the one that participates will be able to earn a higher profit.⁶ If both try to compete in a forward market, then both earn less than if they had waited for the spot market. As is explained below, Allaz and Vila's [3] result has a natural extension to a sequence of auctions.

⁵See [49] for a more extensive discussion.

⁶In particular, the firm participating in the forward market is a *Stackelberg leader* (see [68].

6.4.1 A Simple Two-Auction Example

This subsection considers a very simple example in which the auctioneer wants to sell two objects one at a time to B bidders.⁷ Bidder valuations in each auction are uniformly distributed on [0, 1], that is, the distributions of the bidder valuations in the two auctions are stochastically equivalent. This means that each bidder's valuation for each object in each of the two auctions has the same distribution. This does not mean that any bidder's valuation in one auction is the same as its valuation in any subsequent auction. So, the highest-value losing bidder in the first auction will not necessarily be the highest-value bidder in the second auction. This is in contrast to the assumption of the martingale theorem, that bidder valuations do not change from one auction to the next.

More specifically, it is assumed that each bidder obtains a signal of its value for each object before each auction starts and that the distribution for determining values is the same across auctions (and bidders). This type of model allows for variations where, for example, the lots vary slightly in ways that each bidder might value in a different way, or where bidders' circumstances can shift in a manner which is equally likely to cause value to increase or decrease. In what follows, a few simple examples are used to illustrate how shifts in bidder valuations across auctions can affect whether prices tend to increase or decrease.

A Two-Unit Example

As a base case example, first recall the above example in which there are B bidders competing for two (rather than I) identical objects, and that bidder values do not change at all across the two auctions. It is assumed that each bidder is interested in only one object. For the sake of illustration it will be assumed that bidder valuations are fixed before any auctions start and are drawn from a uniform distribution on [0, 1]. If the objects are sold in a single first-price (i.e., pay-as-bid) auction, then the bidders with the highest two values will win, and the price will be $\frac{B-2}{B+1}$, which is the expected value of

⁷This example follows [10].

the third-highest-value bidder.⁸ On the other hand, suppose the two objects were auctioned one at a time. Then there would be B-1 bidders competing in the second auction, and their expected values would be $\frac{B-1}{B+1}, \frac{B-2}{B+1}, \ldots, \frac{1}{B+1}$. The expected price would be $\frac{B-2}{B+1}$. Therefore, in the first auction, no bidder would want to bid above $\frac{B-2}{B+1}$, and so this would be the expected price in that auction as well. Thus, expected prices are the same whether there is one auction for two objects or two auctions for one object each.

Now, suppose that bidders' values are not necessarily the same for the two objects. This could be the case, for example, if bidders acquire information or their opportunities changes between the two auctions. The latter will occur when, for instance, bidders may engage in other related transactions. Or it can be the case that bidders view the two items as very similar, and bidder valuations for each item are the same except for some random component which is not correlated across items for each bidder. That is, bidder valuations are *stochastically equivalent*.

As above, the auctioneer has two items for sale, and can choose to sell them either all at once or sequentially. The B bidders each can still win at most one item, but now each bidder's valuation for an item is a separate random draw in each auction from the uniform distribution over the unit interval.

In this case, it is easy to compute the expected value of the price if both items are sold at one time in an auction in which both winning bidders pay the third highest offer amount⁹:

$$p^2 = \frac{B-2}{B+1}$$

Then expected auction revenues will be

$$R^2 = \frac{2B - 4}{B + 1}$$

⁸If this were a third-price auction, that is, the highest two bids were to win and pay the third highest bid amount, then each bidder would have a strong incentive to bid its value. The RET states that the expected outcome would be the same as with a pay-as-bid, sealed-bid auction.

 $^{^9\}mathrm{This}$ is the pricing rule equivalent to a second-price rule for an auction with a single object.

If, on the other hand, the items are sold one at a time in two second-price auctions, then the expected price in the second (last) auction will be

$$p_2^{1+1} = \frac{B-2}{B}$$

This means that the expected surplus derived from participating in the second auction will be $\frac{1}{B(B-1)}$. Therefore, (a) a bidder in the first auction will not want to bid up to its full value, as it has a positive expected value from participating in the second auction, and (b) the highest-value loser in the first auction cannot always expect to win the second unit. Thus, the expected value of the price in the first auction will be

$$p_2^{1+1} = \frac{B-1}{B+1} - \frac{1}{B(B-1)}$$

So, expected revenues from selling the items one at a time will be

$$R^{1+1} = \frac{B-1}{B+1} + \frac{1}{B-2} - \frac{1}{(B-1)B}$$

Now $R^{1+1} - R^2 > \frac{1}{B(B+1)}(B-3) > 0$ for B > 3. Therefore, selling the items sequentially raises more revenues when there are more than three bidders.

A Three-Unit, Two-Auction Example

A variant of the second example considers the case with three units and two auction dates. More specifically, suppose the auction originator wants to sell I > 2 units in two auctions. It has the option of selling a fraction x in the first and y in the second. For simplicity, suppose I = 3. In this case, the expected prices in the two auctions depend on whether one unit or two is sold in the first auction.

If one unit is sold in the first auction, then the expected valuations in the second auction of the remaining B-1 bidders are $\frac{B-1}{B}, \frac{B-2}{B}, \ldots, \frac{1}{B}$. This means that the expected price in the second auction will be

$$E(p_2^{12}) = \frac{B-3}{B}$$

Thus, those bidders losing the first auction can anticipate a surplus of $\frac{1.5}{B}$ with a probability of $\frac{2}{B-1}$. So, in the first auction, bidders will shade their bid by

$$s^{12} \equiv \frac{3}{B(B-1)}$$

So the expected price in the first auction will be

$$E(p_1^{12}) = \frac{B-1}{B+1} - \frac{3}{(B-1)B}$$

Similarly, if two units are sold in the first auction, then the expected valuations in the second auction of the remaining B-2 bidders are $\frac{B-2}{B-1}, \frac{B-3}{B-1}, \ldots, \frac{1}{B-1}$. Then the expected price in the second auction will be

$$E(p_2^{21}) = \frac{B-3}{B-1}$$

This means that the expected surplus is

$$s^{21} \equiv \frac{1}{(B-1)(B-2)}$$

So, in the first auction, the expected price will be

$$E(p_1^{21}) = \frac{B-2}{B+1} - \frac{1}{(B-1)(B-2)}$$

Comparing revenues $R^{21} \equiv E(p_1^{21}) + 2E(p_2^{21}) - 2s^{21}$ with $R^{12} \equiv E(p_1^{12}) + 2E(p_2^{12}) - s^{12}$ indicates that $R^{12} - R^{21} = \frac{2}{B+1}[1 - \frac{1}{B}] + s^{12} - 2s^{21}$. This is necessarily positive for large B.

The above two examples illustrate three points.

6.4. STRATEGIC ALLOCATION DECISIONS

- First, the fact that bidder valuations can vary from auction to auction, although they have the same distributions, means that there is an option value from participating in the last auction which does not exist when bidder valuations are constant across auctions.
- Second, the degree of competition will vary across auctions. The firstorder effect means that the larger the number of remaining bidders relative to the remaining auction volume in the last auction, the lower the price in the last auction.
- These examples suggest that when the auction originator has some flexibility about how to divide the total volume available, then decisions about how the volume is divided, and (as will be explained below) the flexibility of the auction originator to shift volume across auctions, can affect average auction prices.

This logic generalizes to some extent to multiple auctions and multiple lots. When bidder valuations across auctions vary, but are stochastically equivalent, the allocation of the auction volume across auctions will affect the expected price in the final auction, and therefore earlier auctions. The relevant order statistic for determining the last auction price is $\frac{B-I+x-1}{B-I}$, where I is the number of items, B the number of bidders, and x the number of lots sold before the final auction. This price will depend on x.

6.4.2 Optimizing the distribution of auction target volume across auctions

This section considers a generalization of the above example with two auctions. It considers a reverse auction. In it the auctioneer wants to *purchase* a total of K units from N bidders, and can choose how to divide this into X units in the first auction and Y = K - X in the second. The question addressed here is what is the optimal way to do so to minimize expected overall procurement costs. The analysis of how to divide a total number of lots to be sold across two auctions is essentially the same.

Given that there are two auctions, and bidders' costs in each auction are, as above, a random, independent, and uniform draw over the unit interval, the second-auction price will be $p_2 = \frac{Y+1}{N-X+1} = \frac{K-X+1}{N-X+1}$. This implies that the average second-period winner will have costs of $\frac{Y+1}{2(N-X+1)}$, and will therefore have an average surplus (for winners) of $\frac{Y+1}{2(N-X+1)}$. This implies a first-period price of $\frac{X+1}{N+1} + \frac{Y}{(N-X)}\frac{Y+1}{2(N-X+1)} = \frac{X+1}{N+1} + \frac{(K-X)}{(N-X)}\frac{K-X+1}{2(N-X+1)}$; the probability that a first-auction loser will be a second-auction winner is $\frac{K-X}{N-X}$. Total costs are then $X\left(\frac{X+1}{N+1} + \frac{K-X+1}{2(N-X+1)}\frac{(K-X)}{(N-X)}\right) + (K-X)\left(\frac{K-X+1}{N-X+1}\right)$. Notice that the first-auction expected price is increasing in the quantity purchased in the second auction. The optimal ex ante value of $X \in \{0, 1, \ldots, K\}$ will minimize

$$C(X,K) = X\left(\frac{X+1}{N+1} + \frac{K-X+1}{2(N-X)(N-X+1)}\right) + (K-X)\left(\frac{K-X+1}{N-X+1}\right)$$
(6.1)

and, if there is no second auction, then the second term in the right-hand side of (6.1) is zero. So, $C(0, K) = C(K, K) = K\left(\frac{K+1}{N+1}\right)$. It will be optimal, ex ante, to divide the auction quantity between the two dates.

Tables 6.1 and 6.2 provide some illustrative calculations. The first example assumes N = 10 and K = 4; the second example considers the case in which N = 400 and K = 250.

x	y	p_1	p_2	Total Costs
0	4	—	.46	1.82
1	3	.25	.4	1.45
2	2	.31	.33	1.29
3	1	.38	.25	1.39
4	0	.46		1.82

Table 6.1: Sequential auctions of 4 lots to 10 bidders

These examples show that there is a benefit in splitting the auction quantity over the two auctions, and some advantage in auctioning more in the first

104



Figure 6.2: 10 Bidders and 4 Lots



Figure 6.3: 400 Bidders and 250 Lots

x	y	p_1	p_2	Total Costs
0	250	0.198	0.63	156.48
100	150	0.377	0.501	112.98
140	110	0.442	0.425	108.60
141	109	0.443	0.423	108.598202
142	108	0.445	0.421	108.598178
143	107	0.446	0.419	108.60
145	105	0.449	0.414	108.63
150	100	0.457	0.402	109.79

106

Table 6.2: Sequential auctions of 400 lots to 250 bidders

auction than in the second, at least when there is a large quantity to be purchased in aggregate. The intuitive reason to divide the auction quantity is that the auction manager gets a larger sample of offers. A bidder at one date is effectively a different bidder at a different date, in that its costs can be higher or lower. Dividing the volume lowers the expected procurement costs, absent strategic bidding. However, bidders will behave strategically. Bidders in one auction will want a higher profit margin than in a later auction. This strategic withholding effect limits how much the auction manager will want to delay procurement.

These examples also assume that the decision to split the auction volume is made before the first auction, and before any bids are received. If the auctioneer can defer the decision to determine the auction volume on each date based on the number of bidders, then it will be possible for the auctioneer to obtain a still better outcome. More specifically, given any volume that the auctioneer might choose to leave for the second auction, it can calculate the expected auction price. This price is what it should use to set reservation prices in the first auction. Such reservation prices in the first auction can only improve its outcome.

6.4.3 Sequential Auctions and Risk

A bidder's preference toward risk can be a factor that enters into a decision about how much to bid in one auction when there are other auctions to

6.4. STRATEGIC ALLOCATION DECISIONS

follow.¹⁰ Consider the case in which a bidder can make a purchase in one of two auctions, each of which is a second-price auction. If the high-value bidder in the first auction lets the second high-value bidder win by bidding low, then in the second auction it can expect to win the object and pay the third highest value. Let $x - V^3$, where V^3 is the (random) value of the third highest value bidder. This will provide the bidder, ex post, with some utility $u(x - V^3)$.

Suppose this bidder is risk averse, and could win the object for the expected value of $V^3 = v^3$. Because of risk aversion, $u(x - v^3) > E[u(x - V^3)]$. As this is the case, the bidder would offer more in the first auction than v^3 to insure against the risk associated with the randomness of the second auction price. Thus, if all bidders are risk averse, then they should each offer a bit more than v^3 to win the first auction, and the expected price should be falling.

If risk aversion is the determining factor in price dynamics across auctions, then prices should decrease. If changing bidder opportunities are, then prices can increase or decrease. Thus, when there are multiple auctions over time, then prices sometimes increasing and sometimes decreasing suggest that risk aversion alone is an inadequate explanation of price patterns. On the other hand, shifts in bidder valuations across auctions would suggest that the price patterns would depend on how the total volume is divided across auctions.

6.4.4 Auction Timing

This subsection now provides another example of how auction timing can affect the auction revenue when bidders learn more about values over time. In particular, it is assumed that there are two bidders and one item for sale. The value each bidder will place on the item is either 0 or 1, and the two bidders have a common ex ante prior, $\alpha \in [0, 1]$, that each value will be high. Moreover, the valuations are independently and identically distributed. So, there are four possible outcomes (0,0), (1,0), (0,1), and (1,1) with probabilities of $(1-\alpha)^2$, $\alpha(1-\alpha)$, $\alpha(1-\alpha)$, and α^2 , respectively. Suppose that the item can be auctioned at one of two dates, early or late.

An early auction means that no ex ante uncertainty is resolved. If the item

 $^{^{10}\}mathrm{This}$ issue was addressed by [52] .

is sold at the early date, then the price, in either a first-price or a secondprice auction, will be α . The winning bidder has zero expected surplus. If, in contrast, the auction is conducted at the late date, then the price will be 1 with probability α^2 . In a second-price auction – or, by revenue equivalence, in a first-price auction – the late-auction price will be 0 with probability $1 - \alpha^2$. Therefore, the expected revenues are α^2 . Notice that the late-auction expected revenues α^2 are less than α , the expected revenues in the early auction.

This result is reversed with a larger number of bidders. If there are *B* bidders, then the probability that no more than one bidder will have a value of 1 is $(1 - \alpha)^B + B\alpha(1 - \alpha)^{B-1}$, and the expected value in a late auction is $1 - [(1 - \alpha)^B + B\alpha(1 - \alpha)^{B-1}]$. The late auction will have a higher or lower expected value depending on *B* and α . When $\alpha = 0.5$ and B = 3, then an early and a late auction have the same expected revenues.

6.5 A Summary of Auction Timing Considerations

This chapter considers situations in which different units of an identical product can be purchased or sold at different dates. This is a fairly common occurrence in many industries, ranging from wine to electricity. The way in which the different units are divided over different auctions at different times can, but need not, affect the outcome. Indeed, a very fundamental result is that strategic arbitrage by bidders will tend to equalize prices across auctions. When this *martingale* result holds, auction design and bidding strategy are both simplified. However, the martingale theorem will not always apply.

When the martingale theorem does not apply, the possibility of shifting purchases and sales over time confronts both bidders and the auction originator with strategic choices. When bidder valuations tend to shift, the above examples suggest there is a tendency of prices to decline in forward auctions, or increase in reverse auctions. However, the auction originator can choose to allocate quantities in a way that can cause prices to either increase or decrease. And competition can shift over time. Therefore, bidders cannot plan on prices falling.
The above suggests that an auction originator's choice of timing will affect the outcome, even without uncertainty. When there is uncertainty, the auction manager will also want to consider the impact of the resolution of uncertainty on competition and the expected prices.

Thus, a firm, government body, or other entity considering auctioning off a set of objects should consider a number of factors in determining whether and how to auction the objects in different events. The relative benefits of different options will depend, in part, on how the allocation affects bidder strategies, as well on how it affects the intensity of competition across auctions. As can be seen above, these choices can affect the outcome in favor of the auctioneer or in favor of the bidders.

110 CHAPTER 6. SEQUENTIAL AUCTIONS OF SUBSTITUTES

Chapter 7

Sequential Auctions of Complements

Summary

This chapter examines auction design and bidding decisions in sequences of auctions in which the outcome of one auction affects a bidder's profits from winning a subsequent auction. At times, a winner of an initial contract or concession will have a stronger incentive to extend its contract or renew its concession in a second auction than it would had it lost the initial auction. This chapter also explains how the bidding strategy in one auction can be affected by the effect the outcome of that auction may have on a follow-up auction for a complementary product or contract. In both cases, we find a tendency toward increasing dominance, that is, the winners of an initial auction have an advantage in follow-up auctions. This also causes bidding to be more competitive in an early auction and less competitive in a subsequent auction.

The complementarities considered can be of two forms. First, the winners of one lot may have a higher value, or a lower cost, for a complementary lot. Second, a winner in one auction may, possibly over time, invest in complementary assets that give it advantages in subsequent auctions.

7.1 Examples of Complementarities

Complementarities are common in auctions. They arise out of decisions of the auction designer about how to divide what is to be purchased or sold into lots. There are many reasons why an auction manager may not want to have one large lot. One primary reason is that different bidders will have different interests. Bidders may be seeking overlapping sets of lots. Or there may be small bidders and large ones. For this reason, an auction may include both large and small lots. This chapter assumes that the lots are sold in sequence; for the most part, it also assumes that the lots are uniform in size.

There are plenty of examples of sequential auctions with complementarities across regions. The European 3G and 4G spectrum auctions were sequential. At least for some of the bidders, the value of a spectrum license in one European Union (EU) country would be positively correlated with having licenses in other, nearby EU countries. There are numerous reasons for this, including the fact that knowledge about the technology and network equipment can be transferred across borders, economies of scale in obtaining in marketing, roaming advantages, and economies of scale in network development and management. In addition, most the 4G auctions included blocks in multiple spectrum bands, which could be either substitutes or complements.

The other type of synergy of concern here is value dependences over time. One fundamental question in many auctions for concessions, licenses, or procurement is the duration of the contract. In many cases, to derive value (in a forward auction) or meet performance requirements (in a reverse auction), the winners must make sunk investments in durable capital projects whose economic life far exceeds the term of the contract, license, or concession. It is often argued that short-duration contracts will not allow the bidder to recover fixed investment costs, or that bidders will need to inflate costs to be able to recover their fixed investments. However, this assertion can only be valid if the winners in an initial auction do not win subsequent renewal auctions. As we shall see, it is often the case that initial winners will win follow-up auctions. Therefore, it is far from obvious what effect an initial contract duration will have on the outcome. Renewal auctions are quite common in some sectors.

One example of this type of problem is spectrum auctions; most spectrum

7.1. EXAMPLES OF COMPLEMENTARITIES

licenses are for a limited term, ranging from ten to twenty years.¹ The question arises as to what happens after the initial license duration or contract ends. While in some cases, as in the FCC auctions, there is a *presumption* of renewal of the contracts, this is far from universally true, and the terms of renewal are often not even described, much less fully specified. It can also take five years or longer to fully build out a license area, and once it is built out, the facilities, such as radio towers, cables connecting towers to the network, and other network equipment, will have value beyond the license term. Moreover, a firm will not want to turn off service to its customers the day its license expires, but it could, in theory, be faced with this prospect.

In the energy sector, in many geographic areas, regional transmission organizations (RTOs) have been organizing bidding procedures for the procurement of capacity. In a capacity auction, the bids are for payments beyond anticipated energy generation service revenue, to ensure capacity gets constructed that otherwise might not be.² The bids often cover a period of no more than ten years despite the fact that the facilities often have useful lives of twenty or more years. These auctions are being conducted both for resource adequacy (that is, to ensure adequate reserve margins) and for renewableenergy procurement. Similar auctions have been planned and conducted for transmission capacity. If the contracts are for ten years but the facility lasts twenty years, then the bidder may have a incentive to ask for nearly twice as much, per unit or per year, for a ten-year contract as it would for a twentyyear one. This logic fails if the first winner is almost certain to win an auction for a contract renewal. However, a shorter contract term does create some risk of losing the contract part way through recovering up-front investments, which is a risk that does not arise with longer contracts.

One auction design option when there are complementarities of the types described above is to permit packages – either prespecified, or allowed to be specified by the bidders. This is often not a practical option. One reason is that package bidding introduces added complexities discussed in Chapter 10. Second, different jurisdictions may not be able to, or may not want

¹The US Federal Communications Commission has been awarding most of the broadband PCS licenses for ten-year terms. See the license-period provisions at www.fcc.gov/wtb/auctions.

²See www.caiso.com for a discussion of the California capacity auction proposal. See www.iso-ne.com/markets for a description of the ISO New England capacity auctions.

to, conduct a simultaneous auction or to create packages. For example, it might be impossible for EU member countries to agree on a single auction for all their 3G or 4G licenses. Third, when the complementarity is over time, the early winners may have strategically different positions than rivals in renewal auctions. The remainder of this chapter considers complementarities in sequential auctions.

7.2 Renewal Auctions

This section examines the influence of contract duration in a sequence of auctions. The section starts with a simple example. It is assumed that there is a franchise auction at some date 0, and a follow-up auction at some other date, 1. Each auction is for a fixed-duration contract – for example, a spectrum license for a fixed number of years. It is assumed that any winner in the first auction must make a sunk durable investment of K > 0 to extract value, but that this investment will not depreciate at all through the end of the second contract period. With the investment of K, a bidder $j = 1, 2, \ldots, N$ can expect to earn $v_j - K$ in the first period if it wins the first auction, or in the second period if it has not already incurred a first-period investment of K; and to earn v_j if it (won the first auction and) has already made the sunk investment of K. In what follows, it is assumed that $2v_1 > K$, that is, there is a positive surplus available.

Further, suppose there are the same $N \ge 2$ bidders in each auction. To simplify the exposition, it is assumed that each auction is a second-price sealed-bid auction. Solving for equilibrium is straightforward.

Start at the last auction. Suppose bidders are indexed so that $v_j \ge v_{j+1}$ for all j = 1, 2, ..., N-1. Also, suppose that $v_1 - v_N < K$, that is, the difference in bidder values is not as large as the sunk costs. Then, if bidder 1 wins the first auction, it will be willing to offer v_1 in the second auction, and the next best offer would be from bidder 2, which would be willing to offer $v_2 - K$. By assumption, $v_2 \le v_1$, so $v_1 > v_2 - K$, and bidder 1 will win the first auction, at a price of $v_2 - K$.

If some other bidder, say bidder h, has won the first auction, then $v_1 - v_h < v_1 - v_N < K$, by assumption, so that bidder h will win the second auction,

at a price of $v_1 - K$. Thus, the bidder that wins the first auction will win the second auction.³ So, the bidder that wins the first auction will also win the second auction, and at a price of $v_1 - K$ if bidder 1 did not win the first auction, or a price of $v_2 - K$ if bidder 1 did win the first auction.

This allows solving for the price, and winner, in the first auction. Bidder 1 would pay up to

$$b_{11} = v_1 - K + [v_1 - (v_2 - K)] = 2v_1 - v_2$$
(7.1)

in the first auction, and the most any other bidder would offer in the first auction, which would be offered by bidder 2, is^4

$$b_{21} = v_2 - K + [v_2 - (v_1 - K)] = 2v_2 - v_1.$$
(7.2)

where b_{ij} is the amount bidder *i* offers in auction *j*. Notice that the term $[v_1 - (v_2 - K)]$ in (7.1) is bidder 1's expected surplus in the second auction. As $v_1 > v_2$ by assumption, $b_{11} > b_{21} = 2v_2 - v_1$. This is an equilibrium in the first auction provided that $2v_2 - v_1 \ge 0$. If $2v_1 - v_2 > 0 > 2v_2 - v_1$, then bidder 1 would win the first auction and have to pay max $\{0, 2v_2 - K\}$, which is the most that bidder 2 could possibly offer. The amount bidder 1 pays in the first auction, $2v_2 - v_1$, exceeeds bidder 1's profits in the first period, $v_1 - K$, whenever $K > 2(v_1 - v_2)$. Bidder 1 is willing to lose money at first because it can recover the losses in the second period. Notice, too, that the price in the second auction is higher (lower) than in the first when $v_2 < (>) v_1 - K$.

This logic extends to the case in which there is a sequence of S auctions, each for a fixed contract duration. In what follows, it is assumed that the fixed and sunk investment is the same as above, and will have an economic life at least as long as that which is spanned by the S auctions. Solving for equilibrium can be accomplished in a manner similar to the above, starting with the last auction.

There are several possibilities in the last auction. Suppose, without loss of generality, that firm 1 is the firm with the highest value of v_j which has won

³If $v_1 > v_N + K$, then bidder 1 could win the second auction after losing the first auction; however, in equilibrium, bidder 1 would win the first auction.

⁴This assumes $2v_2 > v_1$.

a previous auction. Then, firm 1 will have a higher valuation for the last auction, $v_1 > v_j > v_j - K$ for all $j \neq 1$. And the most some other firm will pay is either v_j or $v_j - K$ for some $j \neq 1$. So, the highest-value firm that is a previous winner going into auction S will win auction S and pay $\max_{j\neq 1}\{v_j\}$ over previous winners j, or $\max_j\{v_j - K\}$ if firm 1 won all previous auctions.

Working backwards to auction S-1, we find that the highest-value winner of the previous auction S-2 should win auction S-1 (unless K is not very large and the difference between v_1 and the j which was the highest-value previous winner is relatively large). The willingness to pay of the highestvalue winner in auction S is $v_1 + (v_1 - (v_2 - K)) = 2v_1 - v_2 + K$, assuming 1 won all previous auctions, or $v_1 + (v_1 - v_h) = 2v_1 - v_h$, assuming h is the highest-value previous winner other than firm 1. Both are larger than $2v_h - v_1$, which is the most another bidder would offer in auction S - 1. Inductively, the highest-value winner in auction S - t will win auction S - t - 1, and so all auctions will be won by the same bidder, bidder 1. Bidder 1 will derive a surplus of at least $v_1 - v_2$ in each of auctions $2, 3, \ldots, S$, and $v_1 - v_2 + K$ along the equilibrium path. So, in auction 1, bidder 1 will be willing to offer at least $v_1 + (S-1)(v_1 - v_2) - K$, and should be willing to offer as much as $v_1 + (S-1)(v_1 - \max\{0, v_2 - K\})$. The next-highest-value bidder, 2, should similarly be willing to offer as much as $v_2 + (S-1)(v_2 - \max\{0, v_1 - K\})$. So, the first auction can be for a very high price, larger than $\min\{(S-1)K, (S-1)v_2\}$, with the losses in the first auction offset by profits in the renewal auctions.

Notice that in a sequence of auctions in which the first auction's winner has an advantage in subsequent auctions, the competition in the first auction can be intense. Prices can far exceed the franchise value. However, the losses in the initial contract are offset by the profits from the subsequent renewal contracts. Notice too that the contract duration does not affect overall auction revenues. It does affect auction prices – longer contracts reducing the first-auction price, and not increasing it.

7.3 Sequential Auctions of Complements

This section considers a sequence of auctions for complements. More specifically, it is assumed that the sum of bidder standalone valuations for individual objects is less than the valuations for the packages. Arguably, this is often the case in spectrum auctions, in which bidders realize economies of scale and scope when acquiring more spectrum in a larger geographic area. For example, a firm may place a higher value on a UK license if it also wins licenses in Germany and Italy. Or a firm may place a higher value on a New York license if it also wins a Chicago and a Los Angeles license. There is a large literature on auction design for complements and for package bidding. However, oftentimes it is not possible to conduct auctions for complements at one time, or to allow package bidding.

A very simple example can illustrate how complementarities can affect bidding decisions and prices. Suppose there are two bidders and three lots, and the lots are sold in a sequence of first-price auctions. Further, suppose each bidder derives a value of 1 from two or three lots, and 0 from one lot. The equilibrium price in the first auction of this "chopsticks" example will be 1, and the prices in the second and third auctions will be 0.

Or consider the following straightforward extension of the example from the previous section. Suppose there are two lots, k = 1, 2, and bidders $j = 1, 2, \ldots, n$. Suppose too that for each bidder j, the value of winning lot k is v_{jk} and the value of winning both is $v^j > v_{j1} + v_{j2}$. That is, each bidder places a higher value on the package than the sum of the individual lot values. In what follows, it is also assumed that $v^j > v^{j+1}$ and $v_{jk} > v_{j+1}$ for all j and k.

Consider a sequence of two second-price auctions. To solve for equilibrium, start with the second auction. There, the first auction's winner will have a value of $v^j - v_{j1} > v_{j2}$, and other bidders will have a value of v_{j2} . If bidder 1 wins the first auction, then it will win the second, as by assumption $v^1 - v_{11} > v_{12} > v_{j2}$ for j > 1. If bidder 1 does not win the first auction, the first auction the first auction the first auction, the first auction the first auction the first auction.

So, if 1 wins the first auction, it can expect to win the second auction at a price of v_{22} and derive a surplus of $v^1 - v_{22}$. Therefore, 1 should be willing to bid up to this amount to win the first auction. No other bidder j can derive a greater surplus in the first auction. The most another bidder can expect to win is $v^2 - v_{12}$, as by assumption bidder 2 has the next highest valuations and would at most earn a surplus of $v^2 - v_{12}$ from winning the first auction. Thus, it is a subgame perfect equilibrium for bidder 1 to win both auctions,

the first at a price of $v^2 - v_{12}$ and the second at a price of v_{22} . Note that bidder 1's value minus the total amount paid by bidder $1, v^1 - v^2 + v_{12} - v_{22}$, is positive, as by assumption $v^1 > v^2$ and $v_{12} > v_{22}$.

The above argument will extend to a larger set of complements. More generally, the first auction's winner will tend to have an advantage in auctions for complements. In the above example, the first-auction price can be quite high relative to the second-auction price. For example, suppose $v^2 = Nv = v_{2k}$ for k = 1, 2 in the above example; then the price in the first auction will (N-1)v, and the price in the second auction will be v. Thus, for strong complements, i.e., large N, the price in the first auction will be many times the price in the second.

In the European 3G auctions, there are some complementarities – both on the cost side, in rolling out service, and on the revenue side, at least from roaming revenues. Firms with operations in one country can use the same equipment vendors in another country, and usually benefit from economies of scale. Also, European mobile operators are able to capture a higher share of high-margin international roaming when they have operations in multiple countries. The European 3G auctions were run sequentially (although those in Germany and the Netherlands occurred virtually at the same time). The pattern of prices in the European 3G auctions (Table 6.1) suggests that these incentives were relevant in that early losers were discouraged from competing aggressively in the subsequent auctions.

Country	Date	Per capita revenues
UK	March, April 2000	\$4.90
Germany	July, August 2000	\$4.74
Netherlands	July 2000	\$2.33
Italy	October 2000	\$1.84
Austria	November 2000	0.62
Switzerland	December 2000	\$0.13

Table 7.1: European 3G Auction Prices

One other form of complementarity arises when bidders have budget constraints: the more a bidder spends on one of two or more complements, the less money the bidder will have available for other lots. While, on the surface, the view that budgets can explain the pattern of European 3G auction prices appears plausible, this seems inconsistent with the bidding, unless bidders also were interested in aggregate coverage across Europe. There were 13 bidders that started in the UK 3G auction. The 8 losers did not spend any money, and presumably should have had money for subsequent auctions. In Germany, Debitel (partly owned by Swisscom) dropped out, but Debitel was not involved in the UK. So, the 8 losers still should have had money to bid in the Netherlands, Italy, Austria, and Switzerland, as well as in the sealed-bid tender in France (where one of four 3G licenses went unclaimed at a minimum of 2 billion euros).

7.4 Sequential Intellectual Property Rights Auctions

Another area in which sequential auctions occur is for new technology. Over time, technology improves. Often, the firms developing the technology are not the ones that use it directly; rather, they license it to downstream competitors. Those owning the intellectual property (IP) can conduct competitive bidding processes to determine which firms will be able to use the technology. Often, as is assumed in this section, the IP owner will prefer to license a technology to one firm, rather than to multiple firms. This will be the maintained assumption in this section. The question addressed is whether the technology leader will tend to have an incentive to outbid rivals with inferior technologies. In other words, if a firm gains access to technology that is better (e.g., one generation ahead), possibly through an auction, then will it win an auction for rights to the IP for a next generation technology that represents an improvement over previous versions?

To capture these ideas, it is assumed that there are two firms, and each has access to a technology characterized by a single cost parameter. More specifically, it is assumed that a firm with generation n technology will have constant unit costs, c^n , and that $c^n > c^{n+1}$ for all n. It is assumed that there are no other costs.⁵

⁵Much of this follows [62]. See also [72], and [27].

First consider a simple two-period model. Each firms starts with a basic technology, with cost parameter c^0 , available to any firm. Before each period there is an auction for the new technology. The first auction is for technology with costs c^1 and the second for technology $c^2 < c^1$. It is also assumed that technology 1 can be used for the second period before the auction for technology 2 is held, and that technology 2 will have a value derived only over the second period.

What follows derives the incentives of the technology leader to win the auction for technology 2. These incentives depend on the nature of the competition between the two firms. First consider the case in which the firms compete in price, and there are perfect substitutes. If the leader wins the second auction, it will earn

$$\Pi^{L} = \max \left[D(p)(p-c^{2}) | p \le c^{0} \right]$$

where D(p) represents market demand at price p. On the other hand, the lagging firm will earn

$$\Pi^{l} = \max \left[D(p)(p-c^{2}) | p \le c^{1} \right]$$

Therefore, $\Pi^l < \Pi^L$, and so the leading firm will tend to win the second auction.

Now consider the case in which there is a follow-up innovation that would achieve costs $c^3 < c^2$. Then the incentives of the lead firm to win are

$$\Pi^{L} = \max \left[D(p)(p-c^{3}) | p \le c^{0} \right]$$

and the lagging firm would offer as much as

$$\Pi^L = \max \left[D(p)(p-c^3) | p \le c^2 \right]$$

in the second auction.⁶ If there is a sequence of n non-drastic innovations,⁷ i.e.,

$$\operatorname{argmax}\left[D(p)(p-c^n)|p \le c^0\right] = c^0,$$

then a similar argument implies that the same firm would win each auction. It is possible to work backwards and calculate the prices in each auction. To simplify exposition, it is assumed that D(p) = D for all $p \leq \overline{p}$.

The price in the last auction would be $[c^{n-1} - c^n]D$. In the next to last auction, the price would be

$$D\left\{ [c^{n-2} - c^{n-1}] + [c^{n-2} - c^n] - [c^{n-1} - c^n] \right\} = 2D[c^{n-2} - c^{n-1}]$$

A similar calculation can be used to establish the price of the *m*th-fromlast auction as $mD[c^{n-m} - c^{n-m+1}]$. So, with equal cost improvements, i.e., $c^k - c^{k+1} = c^j - c^{j+1}$ for all j, k, the low-cost firm will always win, and prices will be decreasing from one auction to the next.

This result has been called an *increasing dominance* result.⁸ It is sensitive to the nature of the competition in the market for products or services using the technology.

To see this, consider the case in which the firms are Cournot competitors with constant unit costs c that depend on the vintage of the technology, v, that each firm has. Suppose that one firm has a base technology c^0 and the other has a newer technology, $c^1 < c^0$, and that there is an auction for an improved technology $c^2 < c^1$. Suppose, too, that there is linear demand P = a - bQ, where Q is the total output of the two firms. Straightforward calculation shows that when the leading firm wins the auction for technology c^2 , and the other firm has technology c^0 , then the Cournot equilibrium quantities will be

$$q^L = \frac{a - 2c^2 + c^0}{3b}$$

⁶A similar argument implies that the lead firm would still win if it had costs c^2 and the other firm had costs c^1 .

 $^{^7\}mathrm{Here}$ drastic means an innovation so large as to make the old technology economically irrelevant in the market.

⁸See [72], or [27].

for the leading firm, and

$$q^l = \frac{a - 2c^0 + c^2}{3b}$$

for the lagging firm. Profits will be $\frac{(a-2c^2+c^0)^2}{9b}$ for the leading firm and $\frac{(a-2c^0+c^2)^2}{9b}$ for the lagging firm. On the other hand, should the lagging firm win the auction for technology c^2 , profits would be $\frac{(a-2c^2+c^1)^2}{9b}$ for the firm that was previously the lagging firm, and $\frac{(a-2c^1+c^2)^2}{9b}$ for the other firm. So, the firm that starts with technology c^1 will want to offer

$$L = \frac{(a - 2c^2 + c^0)^2}{9b} - \frac{(a - 2c^1 + c^2)^2}{9b}$$

and the other firm will offer

$$l = \frac{(a - 2c^2 + c^1)^2}{9b} - \frac{(a - 2c^0 + c^2)^2}{9b}$$

to win technology c^2 . Here L can be larger or smaller than l, depending on the values of c^0 , c^1 , and c^2 . For example, if a = 10 and $b = \frac{1}{9}$ and if $c^0 = 3$, $c^1 = 2$, and $c^2 = 1$, then l > L. But L > l if $c^0 = 5$, $c^1 = 4$, and $c^2 = 3$.

7.5 Auctions for Complementary Inputs with Imperfect Competition in Downstream Markets

The above analysis also is useful in examining what happens when there is an auction for an upstream input that can affect downstream competition. Spectrum auctions are one example of this. Most spectrum auctions now are largely for incumbents competing to expand capacity and lower costs. The technology of spectrum is such that if one firm has more spectrum than its rival, it can offer service at a lower cost per unit of capacity – and thus, with higher bandwidth and better quality.

To model this, suppose there are two firms that currently have costs c_j^0 , j = 1, 2, and that there is an auction for a single block of spectrum that will

reduce their marginal cost to $c^1 < \min\{c_1^0, c_2^0\}$. It is also assumed that, as in the previous subsection, the firms engage in Cournot competition, so that the price P = a - bQ, where Q is total output of the two firms.

The incremental profit for a firm that wins the additional block is

$$\Delta \Pi_i = \frac{(a - 2c^1 + c_j^0)^2}{9b} - \frac{(a - 2c_i^0 + c^1)^2}{9b}$$

Whether the firm with lower initial costs will be willing to outbid its rival will depend on whether its incremental profit will be higher. With Cournot competition in the downstream market, it is possible that either the leading or the lagging firm can win, as in the previous subsection. On the other hand, if there is Edgeworth-Bertrand competition, the same firm will win each auction.

This issue is important in setting policy, because spectrum is regulated by government agencies that often set caps to ensure that the market does not become excessively concentrated. These types of models suggest that the market for wireless services can, under some conditions, become increasingly concentrated if there is a sequence of auctions that gradually release and reallocate spectrum over time. This section has merely served to introduce a framework for assessing the issue, rather than try to provide a complete analysis.

This section also serves to show that increasing dominance, and therefore concern about spectrum caps, can – in some, but not all, circumstances – be justified. The analysis of this issue is incomplete, and will be addressed in more detail in Chapter 10.

The above only considers sequential auctions. A regulator may have the option of selling blocks sequentially or in a single auction, and of allowing a single auction to include many small blocks or only a few (possibly only one) large ones. A simultaneous auction of all the spectrum blocks, or of multiple licenses for new technology, can have different outcomes, including the possibility of leapfrogging.

To see that leapfrogging may occur, consider the case in which there is a single large block of spectrum available for bid, possibly instead of several small blocks. Also, suppose there are two firms bidding: a leader, L, and a lagging, smaller firm, s. Let $\Delta \pi_1^L > \Delta \pi_1^s > 0$ denote the incremental profits of the two firms. If the leading firm knows $\Delta \pi_1^s$, it will bid just a bit more than this and become more dominant. However, if the leading firm does not know $\Delta \pi_1^s$, then it will wish to maximize $[\Delta \pi_1^L - \Delta \pi_0^L - b^L] \times F(b^L)$, where $\Delta \pi_0^L$ is the leading firm's profit if its rival wins, b^L is what it bids, and $F(b^L)$ is the probability that this bid wins. This results in a first-order condition that

$$\Delta \pi_1^L - \Delta \pi_0^L - b^L = \frac{F(b^L)}{F'(b^L)}$$

This necessarily means that leapfrogging can occur.

7.6 Summary

This chapter has considers sequential auctions for complements. One striking result, which appears quite robust, is that the winner of a first auction for one of two or more complements does not really face an exposure problem. Admittedly, these models are highly simplified and abstract, and bidders do face risks that are not captured by them. However, the initial winner in the first auction will tend to have an strong advantage, which will tend to increase in a sequence of auctions for complements. The fact that an initial winner can lose a subsequent auction does not totally offset the benefits of winning an initial auction, but may somewhat reduce incentives to bid beyond values initially in order to win the market.

A second striking result is that prices will not be uniform. The first auction will tend to attract prices well beyond values, and subsequent auctions will lack intense competition. So, for example, the initial award of a franchise or license can be very intense, but the re-auction, or renewal, is then likely to see much less or almost no competition. Prices in any one auction are not linked to underlying values for that auction, but to the entire sequence of auctions. These issues arise in many auctions. For example, many auctions of licenses or franchises have limited durations, and do not necessarily specify all the renewal provisions in advance. The analysis here suggests that license durations that exceeds the life of the investment may not be essential. However, the analysis provides a way of assessing the important role that the renewal provisions can take in determining bidder incentives and the outcome in the first of a series of auctions.

126 CHAPTER 7. SEQUENTIAL AUCTIONS OF COMPLEMENTS

Chapter 8

Single-Product Auctions

M any auctions include identical lots of a single product. Rather than auction each lot individually, an auctioneer can ask bidders to submit demand (or supply) curves to express how much each bidder wants at any given price. An alternative way for an auctioneer to elicit such demand (or supply) is in the form of a clock auction, in which the auctioneer names a monotone sequence of prices and asks bidders to indicate how much they want at each price. This chapter describes such auctions, and shows when a demand-function (or supply-function) game is strategically equivalent to a clock auction. It derives equilibria for demand- and supply-function games with and without uncertainty. It also explores the effect of capacity constraints. Lastly, it describes a number of applications of demand- and supply-function auctions and clock auctions.

8.1 Introduction

Often, auctions include many units or lots of an identical product. Indeed, virtually all auctions in the energy sector allow bidders to submit offers for one or more lots or blocks. The blocks can be 25 MW entitlements, rights to 100 MW or 1 MW of transmission rights, or other products. Multi-unit auctions are found in many other sectors. For example, the agricultural cooperatives Fonterra and Ocean Spray each run monthly auctions: In Fonterra's

auctions bidders can buy any number of full container loads of various forms of milk products, and Ocean Spray's auctions are for cranberries and related products.

An auctioneer wishing to sell, or purchase, many units of a single product will need to decide whether to include all units in one bid, in a single auction, or in separate auctions. If multiple or all units are auctioned at one time, then the auctioneer needs to decide on what form bids should take and how to elicit information from bidders about how much they want at different prices. Bids can be demand functions, i.e., separate prices for each unit. There can be restrictions on the form of demand function. The auction can allow bidders to submit one price for all units or different prices depending on the quantity bid. Alternatively, as in a clock auction described below, the auctioneer can announce a sequence of increasing prices and, at each price, ask each bidder to indicate its demand.¹ The auctioneer may also need to consider the potential impact of strategic withholding by bidders. This chapter provides a review of some of the more commonly used multi-unit auctions and describes their properties.

There are several ways in which multi-unit auctions are typically conducted. Probably the most common method for auctioning identical units of a product is by means of sequential auctions, and one at a time. This type of auction is easy to organize. But, as what follows should make clear, this may not always be the best method. Multiple auctions for each available unit can be held in parallel when the auctions are conducted online.

At times the auction format is influenced by indivisibilities in underlying production capacity. For example, in the energy sector, supplies are typically linked to generating units that have fixed capacities. Possibly for this reason, many reverse energy auctions allow bidders to submit supply functions, and to choose the form of the bid function is based on generation facility capacities. The auctioneer then aggregates the supply-function bids, and determines the price and the allocation based on market clearing of the aggregate supply and requested demand.

The *clock auction* is another type of multi-unit auction in which all units are sold at one time. As is explained below, clock auctions are in some cases

 $^{^1\}mathrm{In}$ a reverse auction, the price sequence would be decreasing and the offers would be supplies at each price.

strategically equivalent to supply-function auctions. This strategic equivalence fails to hold when bidders are provided information about rival bids between clock auction rounds. The combinatorial clock auction (CCA) is a recently developed variant of the clock auction in which bidders can submit any number of package bids in a final sealed-bid round, and prices are based on a modified second-price rule.

This chapter examines auctions for multiple identical units of some product. The next section describes clock auctions. Clock auctions can differ in the information provided to bidders between rounds, and in the pricing rule. Clock auctions are compared with sequential auctions. The next section describes demand- and supply-function auctions under complete information. Then, the outcome of clock auctions and demand-function auctions are compared. In all these multi-unit auctions, bidders' strategic withholding can adversely affect the outcome.

8.2 Simultaneous vs Sequential Auctions

The analysis of the last chapter indicates that the expected price in a sequence of auctions for identical products is a *martingale*, that is, the expected price in one auction equals the actual price in the previous auction, provided bidders' valuations are the same across auctions at different times. This result no longer holds when bidder valuations can change. Should bidder valuations be static, then, as explained in the previous chapter, the price and allocation are independent of how many units are sold in each auction. Indeed, the highest-value bidders win, and all winners pay the same expected price.

This is no longer the case when bidder valuations can change. In that case, price can vary from unit to unit across auctions. What this means is that a higher-value bidder can lose in an early auction to a lower-value bidder, and fail to win later. This type of outcome would come about due to a possible miscalculation by the losing higher-value bidder or to an unexpectedly high level of competition in the later auctions. Unexpected fluctuations in demand across auctions can be difficult to plan for. However, even with little or no fluctuation, high-value bidders can lose to lower-value ones in sequential auctions. This type of outcome reduces the average bidder surplus as well as overall auction revenues. Thus, sequential auctions require bidders to make complex calculations and forecast future demands. As is explained below, in simultaneous clock auctions such guesses are unnecessary, eliminating the potential for mistakes.

However, sequential auctions do have advantages when demand across bidders can shift over time. As explained in the previous chapter, overall auctioneer revenues (for a forward auction) will tend to be greater when the auction volume is split across two or more auctions. If this benefit is small, then the possible inefficiencies due to bidder miscalculations and bidding errors can offset the benefits of sequential auctions.

Clock Auctions 8.3

A clock auction is probably the simplest multi-unit auction. In a clock auction the price starts below market clearing levels, and the auctioneer raises prices in a series of small discrete steps. After each upward tick of price, bidders indicate how many units they want to purchase. Prices keep rising until there is no longer any excess demand. When the clock stops, the remaining bidders are the winners, and each bidder is usually asked to pay the last posted price. Other pricing rules and stopping rules can be used. For example, all bidders can be required to pay the last price for which there is excess demand. Another alternative is that the clock price increases until there is no longer *any* demand, and each bidder can be required to pay the price at which it drops out.

The clock auction is essentially a simple adaptation of the classic Walrasian tatonnement process.² The main distinction between the clock auction and the Walrasian auction is that the supply side is absolutely inelastic in the clock auction, but can respond to price in the Walrasian auction. The set of clock round offers trace out a demand schedule. The auction thus ends when the price rises to the point where demand drops to or below supply. In the remainder of this section, I briefly describe a few properties of the clock

 $^{^{2}}$ See [74].

8.3. CLOCK AUCTIONS

auction and a few variants.³

The simplest case is where each bidder will purchase at most one unit. In this case, each bidder should remain in the auction until the price reaches its value. If a bidder were to drop out before the price reached its value, it would derive no benefit from doing so – and this can only save rivals money. There is similarly no point in waiting until after price goes past value, as the bidder may then overpay. The bid strategy is essentially the same as in a standard English auction.

If bidders pay the last price for which there is excess demand, then it is a dominant strategy to bid up to value and then drop out. In other words, bidders will want to bid value, independently of rivals' values and bids. On the other hand, if bidders pay the last price posted, which means the first price posted for which there is no longer excess demand, then it is possible for a bidder's offer to affect how much it pays. In this case, a bidder's strategy can depend on the tie-breaking and rollback provisions. Specifically, if at the next to last price demand exceeds supply, but at the last price there is excess supply (i.e., there is overshooting), then a bidder may want to drop out a bit before the price reaches its value. This is especially true if the increments are large or the bidder believes that overshooting is likely.

Thus, the rule that determines prices can have some effect on the bidders' strategy. If bidders pay the highest losing offer, then bidding value is a dominant strategy. If bidders may have to pay the lowest winning bid, then they can have incentives to drop out early. Note that these differences are likely to be small if there are a large number of bidders and values are closely bunched, that is, if the gap between the highest losing offer and the lowest winning offer is small.

8.3.1 Clock Auctions and Sealed-Bid Auctions

The decision facing a bidder that is seeking to purchase up to one unit in a clock auction is very similar to the decision facing a bidder in a sealed-bid auction. Indeed, at times the decisions will be strategically equivalent. When this is the case, the outcomes of the two types of auctions will be the same.

³Some of what follows is in [54].

The only differences between a clock auction and a sealed-bid auction are in the information available to bidders when having to decide on crucial bids. In the extreme case of a clock auction in which the auctioneer (a) never reports any information about excess demand and (b) increases price until there is zero *demand*, and not just zero *excess* demand, the decision facing a bidder is identical in a second-price sealed-bid auction and in a clock auction in which the bidder pays the high losing bid. A bidder has no information about rivals' participation in the clock auction. And so, each bidder is just choosing a price at which to drop out of the auction. This stopping price should equal the bidder's value, and be exactly the same as in the secondprice sealed-bid auction.

Similarly, the strategic decision facing a bidder in a clock auction is the same as in a sealed-bid auction when (a) the clock auction price is the lowest winning bid and (b) the bidder pays its bid amount in the sealed-bid auction. In both cases, the bidder will want to shade its bid, that is, drop out below value. The less it bids, the greater is its surplus, but the lower the chance of winning. And, as in the clock auction, the bidder gets absolutely no information about rivals, so its choice of when to drop out is just based on when it thinks the next-highest-value rival would drop out.

This strategic equivalence of clock auctions and sealed-bid auctions breaks down when information is reported to bidders between price increments (rounds). When, in a clock auction, bidders pay the highest losing bid, information about rival demand does not affect strategy except to the extent that this information affects bidder valuations. This means that bidders having affiliated values will tend to bid more aggressively in the clock auction than in the sealed-bid auction.

If the clock auction price is set at the lowest winning bid, then the strategic equivalence of sealed-bid auctions and clock auctions breaks down, even with independent private values. When bidders know that there is no excess demand, they can stop bidding. Thus, in a clock auction, bidders won't always need to bid up to the point where the benefit of increasing the probability of winning by offering more just offsets the reduction in surplus. Often, the auction will end sooner. Thus, there can be some gain to having the additional information in the clock auction. It is possible that a bidder will choose the same price at which to stop bidding in a clock auction as it would offer in a first-price auction, but this would depend on the distribution of bidder valuations. Expected revenues would be less, though, as in some cases the winning bidder would not have to pay as much as it would offer in a first-price auction.

8.3.2 Clock Auctions, Sealed-Bid Auctions, and Market Power

We now consider the case in which a bidder may want to purchase more than one unit. In this case, its decision to reduce demand can reduce the price it pays for the units it does purchase. In a *second-price* auction, that is, an auction in which winning bidders all pay the highest losing bid amount, it is a dominant strategy for each bidder to bid its value. However, if some bidders are seeking to purchase two or more units, each can have an incentive to bid less than its value in some cases.

A simple example can illustrate the incentives to withhold demand in this case. Suppose, to start, that there are three bidders, each having values drawn from a uniform distribution on [0, 1], and that there are two units for sale. One bidder has a demand for two lots, and the other two bidders only want one lot each. The top two bidders would win and pay the third highest bid amount. The expected values in rank order are $\frac{4}{5}$, $\frac{3}{5}$, $\frac{2}{5}$, and $\frac{1}{5}$. Assuming the large bidder does not shade its bid, the expected price would be $\frac{2}{5}$.

Now consider the incentives of the bidder seeking two units to shade its bid. The possible rank orders (where L denotes the large, two-unit bidder, and s one of the two small, one-unit bidders) are LLss (i.e., the two-unit bidder L has the highest two valuations, and the other two bidders s have the lowest valuations), LsLs (i.e., L has the highest and third highest valuations), LsLs, sLs, and ssLL.

In the case LLss, the two-unit bidder will win two blocks with an expected price of $\frac{2}{5}$. This occurs with probability $\frac{1}{12}$. So, this bidder, by bidding 0 for its second unit, loses, and loses an expected payoff of $\frac{1}{5} \times \frac{1}{12}$. On the other hand, bidding 0 for the second unit will allow this bidder to pay on average $\frac{1}{5}$ instead of $\frac{2}{5}$ in the cases LsLs, which occurs with probability $\frac{1}{6}$, and sLLs, which occurs with probability $\frac{1}{12}$. In the other three cases, underbidding does not affect this bidder's expected payoff. So, this bidder gains an expected

net return of $\frac{1}{6} \times \frac{1}{5} = \frac{1}{30}$ by strategic underbidding. This type of strategic demand reduction has been noted as possible both in spectrum auctions⁴ and in the energy sector.⁵

Note that if bidders pay the lowest winning bid, then the outcome will generally be the same as in the case where bidders pay the highest losing bid, assuming each bidder is seeking to win only one unit. When bidders are seeking multiple units, they will also have some incentive to underbid.

8.4 Supply-Function Auctions

This section considers reverse auctions in which bids are in the form of supply schedules. More specifically, each bidder informs the auctioneer how much it is willing to supply as a function of the price it is paid. Thus, the bidder may offer a quantity q_0 at price p_0 and a quantity $q_1 \ge q_0$ at price $p_1 > p_0$. These supply-function bids are a common way for power to be offered in energy auctions.⁶ The reason for the use of supply functions is on the technical side: the offers can be matched to the characteristics of the plants providing energy⁷ (although, in many cases, energy offers are in the form of financial contracts – the supplier guarantees a spot price determined by a grid operator). As Green and Newbery [32] noted, the reliance on supplyfunction bids may have the unintended consequences of resulting in higher (i.e., less competitive) transaction prices.

A very simple example of a supply-function game is one in which two plant owners, say having $M = M_1 = M_2$ MW of capacity, can offer to sell it in to the grid operator a price p_j for operator j = 1, 2. Suppose too that the needed supply, S, satisfies $\frac{M}{2} < S < M$. For example let S = 100 and M = 60. Also assume that the grid operator, who serves as the auctioneer, will cover the fuel costs, and so bids are just for rights to use the capacity. The grid operator is assumed to have a reserve price R, which for purposes of this example we take to be 100. It is also assumed that operators receive their bid price.

 $^{{}^{4}}See [18] and [34] and [18] .$

 $^{{}^{5}}See [36], [75] and [13]$.

 $^{^{6}}$ See [32], [9], and [48].

⁷See [9] for a discussion.

8.4. SUPPLY-FUNCTION AUCTIONS

It is clear that no operator would ever accept a price less than $\frac{S-M}{M} \times R$, or $66\frac{2}{3}$ in this example, the reason being that the bidder can always sell S-M at the reservation price. However, when one bidder chooses a price strictly above $\frac{S-M}{M} \times R$, its rival will have incentives to undercut. Thus, there is no pure strategy equilibrium. It turns out that choosing a probability distribution function F(p) such that each rival earns the same expected profit, independent of the price the rivals choose between $\frac{S-M}{M} \times R$ and R, is a mixed strategy equilibrium. This is the case if

$$Mp[1 - F(p)] + (S - M)pF(p) = k$$

or

$$F(p) = \frac{M - \frac{k}{p}}{2m - S}.$$

It can also been seen that $k = M \left[\frac{S-M}{M} \times R\right]$. In the numerical example, $F(p) = \frac{60 - \frac{4000}{p}}{20}$ for $66.67 \le p \le 100$.

More generally, consider the case of two firms bidding into a market (auctioneer) demand function Q = D(p).⁸ Let $S^{j}(p)$ denote the supply schedule offered by firm j = 1, 2. What follows is a description of some of the necessary conditions for an equilibrium set of offers $S^{*1}(p), S^{*2}(p)$ and resulting equilibrium price p^* and quantity Q^* . The assumption that the auctioneer incurs all variable costs, as is often the case in energy markets, is maintained.

Note if there is an equilibrium, unique or otherwise, then $D(p^*) = S^{*1}(p^*) + S^{*2}(p^*)$, that is, supply must equal demand. If there is no equilibrium, it is assumed that neither firm earns any profit. Now, consider *any* pair of firm quantities (q^1, q^2) , such that the price p satisfies

$$p = D^{-1}(q^1 + q^2)$$

What follows derives supply functions for the each of the two firms that can support these prices as an equilibrium outcome. Notice that given $S^i(p)$, firm j will want to maximize $p[D(p) - S^i(p)]$ with respect to p. This implies that the following first-order conditions need to be satisfied:

⁸This section is loosely based on [43].

$$[D(p) - S^{i}(p)] + p[D'(p) - S^{i'}(p)] = 0$$
(8.1)

But (8.1) is satisfied whenever⁹

$$S^{i'}(p) = D'(p) + \frac{D(p) - S^{i}(p)}{p} = D'(p) + \frac{q^{j}}{p}$$
(8.2)

However, multiple values of the price will satisfy (8.2). Thus, there will be multiple equilibria. For example, if demand is linear [D(p) = A - Bp], then linear supply schedules $S(p) = s \times p$ will satisfy (8.2) whenever $s = \frac{A}{p} - 2B$.¹⁰

Suppose there is uncertainty, e.g., suppose $D(p) = A - Bp + \epsilon$, where ϵ is a nondegenerate random variable, that is, the demand has a random component. In this case, the supply schedule each bidder chooses has to satisfy the first-order conditions in (8.1) for any realization of the random variable. This means that a pair of supply functions that are an equilibrium for one particular realization of demand need not still be one for others. It can be shown that there will be a unique pair of supply functions that will be an equilibrium for all realizations of demand.

8.5 Conclusions and Applications

This chapter has reviewed clock auctions and supply-function auctions. The two are very similar, and, as explained in Section 3, are identical when there is no information reported to bidders until bids go to zero in the clock auction. When bidders see when rivals drop out, or have some aggregate measure of rival offers, then the clock auction allows bidders to bid a bit more aggressively than in a sealed-bid supply-function auction.

In practice supply-function (or demand-function) auctions are common in some sectors, most notably the energy sector. There are often restrictions

⁹If marginal costs are positive, the first-order condition becomes $[D(p) - S^i(p) + (p - C'((D(p) - S^i(p)))][D'(p) - S^{i'}(p)] = 0.$

¹⁰As long as demand is concave and costs are convex increasing, the second-order conditions will also be satisfied.

on the form of bids, e.g., bidders can submit only step function bids or piecewise linear schedules. These restrictions tend to reduce the complexity of the bidding process.

A key issue is how well clock and supply-function auctions work. Note that bidders can have market power in these auctions. This is most clearly seen in the simple capacity example of Section 8.3.2, in which each of two bidders offers its capacity to a grid operator. This result can be generalized; more specifically, it can be shown that in cases in which bidders can be pivotal, bidders' offers will always be bounded away from their competitive levels. This possibility was noted in [32], seen in practice, and proven in theory.¹¹

The logical question is what are the alternatives to these type of auctions. One of the key issues in auction performance in multi-unit auctions is the incentives of bidders to make offers close to marginal values. The auctioneer can choose to reduce volume if offers are insufficiently competitive, and this can serve to elicit more competitive auctions from bidders.¹² However, this may only delay the inevitable. If bidders know that the auctioneer has to conduct another auction for what is left over, they will want to wait for the re-auction. If the conditions of Weber's martingale theorem hold, that is, bidders' values are static, then the re-auction will not help. On the other hand, if the bidders' values can increase or decrease, then, as explained in Chapter 6, the auctioneer would be well served by reserving the option of deferring some purchases or sales to a later auction, and making this decision endogenous, that is, based on the bids.

¹¹See [48].

138

Chapter 9

Simultaneous Auctions

A novel approach to auctioning multiple objects began when the United States Federal Communications Commission (FCC) conducted the first simultaneous ascending multiple-round (SMR) auction in 1994. This approach allowed bidders to link their offers for different objects – they could arbitrage substitutes and limit the risk of winning one of a set of complementary objects in SMR auctions. In other words, the SMR auction was probably the first type of auction that allowed bidders to make simultaneous decisions for a set of objects whose values were related. Prior to this, these decisions had to be made independently, creating all sorts of difficult decision problems for bidders. The technology for auctioning multiple objects with interdependent values is still improving. This chapter provides an introduction to the theory and practice of multi-object auctions.

9.1 The SMR Auction

It might not be hyperbole to state that the adoption of the simultaneous multiple-round (SMR) auction by the US Federal Communications Commission (FCC) was the single most significant event in the application of game theory to address practical decision problems. This auction format was developed de novo to help the FCC, a governmental agency, more efficiently allocate an extraordinarily valuable resource, generating over \$10 billion in nominal auction revenues in the first few years. The problem facing the FCC was to develop an auction mechanism that would efficiently allocate spectrum licenses, that is, licenses to use specific frequency channels over different geographic areas. The licenses could be substitutes, as where there are two or more similar or identical blocks covering the same geographic areas, or complements, as where a bidder is seeking different bands covering the same area for different needs, or in different areas (where coverage is an issue).

This allocation presented several challenges. One is that in parallel or sequential auctions substitutes for identical products often sell for different prices. In such cases a losing bidder for one lot can have a higher willingness to pay than a winning bidder for an identical lot. This also means that the outcome of the auctions is necessarily inefficient. Reallocation through aftermarket transactions cannot be relied upon to correct for such inefficiencies unless transaction costs are low. This same logic extends to close but imperfect substitutes; it is possible that a higher-value lot will sell for less than a lower-value one.

Another challenge is the *exposure* problem: a bidder entering separate bids for two or more complements may end up overpaying if it wins part, but not all, of its desired package. For example, if one block is worth 10, but two blocks are worth 50, a bidder that offers 20 per block may win only one. Allowing bidders to submit all-or-nothing package bids is one approach to this exposure problem. However, package bidding introduces other complexities, as is explained in more detail in Chapter 10. And auction rules that help one bidder address the exposure problem can create a bias in favor of a large package bidder over bidders seeking smaller packages.

The chapter starts with a brief description of the SMR and clock auction formats. A simultaneous clock auction differs from an SMR auction in that bidders name only quantities based on prices set by the auctioneer in each round. This chapter reviews properties of the SMR auction format. Then it provides a review of the experience with SMR and simultaneous clock auctions. In particular, it presents both good and bad experience with the SMR auction. Finally, it explains how the performance of an SMR auction (that is, how close it comes to achieving an efficient outcome) and the revenue raised can be surprisingly sensitive to seemingly minor decisions about auction configuration and parameter settings. Some guidance is provided about how to ensure that an auction outcome is more likely to be efficient, and to counter bidder efforts to exert market power.

9.2 Introduction to the SMR Auction

9.2.1 A Brief History of the SMR and Clock Auctions

In 1994 the US FCC faced a Congressional mandate to allocate and assign various combinations of regional and national spectrum licenses using auctions.¹ Because of incumbent users on the spectrum, the new licenses would, in some cases, be awarded to firms competing with existing licensees, and in other cases, existing licensees would be seeking to add to their frequency holdings. In addition, both bandwidth and geographic efficiencies meant that in some instances a pair of licenses covering the same or adjoining regions could be substitutes or complements. The FCC's mandate included efficiency and the promotion of deployment of advanced telecommunications services as primary goals, with revenue a secondary goal.²

Thus, the FCC faced a difficult auction design problem – it wanted an auction that would efficiently allocate licenses that could be complements or substitutes. This meant that the auction should allow efficient arbitrage in the case of substitutes. And, in the case of complements, the auction should allow bidders to effectively manage the exposure problem when bidding for a package of complementary licenses. Moreover, the FCC would not want to allow bidders to win too much spectrum so as to create market power. An additional complication was that there were different options for dividing the spectrum bandwidth with respect to geographic partitioning of the US and its territories.

¹See http://wireless.fcc.gov/auctions/ and [47]. Note that for spectrum auctions, the term *allocate* is usually reserved for the type of use of the spectrum (television, mobile communications, public safety, etc.), and the term *assignment* is used for the process whereby individuals or firms are awarded rights to use the spectrum.

 $^{^{2}}$ See [47].

9.2.2 SMR Auction and its Variants

This section provides a very brief description of the basic SMR and simultaneous clock auction rules. As noted above, the first SMR auctions, introduced by the US Federal Communications Commission, were initially introduced for the sale of a number of lots, some of which could be substitutes and some could be complements in a single auction. In any SMR or clock auction there are multiple units or blocks available in a single auction. And bidding occurs in a sequence of rounds.

Prior to the start of an SMR or clock auction, the auctioneer posts some application and eligibility requirements. One component of the initial eligibility requirements includes the application for eligibility points, either based on the scale of activity or set by the auctioneer for the lots available. Bidders must include in their application some indication of how many points they want to be able to bid for. A bidder is never allowed to win more points than it indicated in its application. This is a component of auction rules designed to encourage increasingly serious bids as the auction progresses.

In the first round of any SMR or clock auction, a bidder is allowed to enter bids for any combination of lots whose total activity points do not exceed that bidder's initial eligibility. If a bidder fails to meet its *activity target*, which is some fraction, up to 100%, of its eligibility, it will see its eligibility decrease in the next round. For example, if a bidder's eligibility is 200 activity points, the required activity percentage is 80%, and this bidder is active on 64 points in the first round, its eligibility for round 2 will fall to $80 = \frac{64}{80\%}$. And once a bidder loses eligibility, it can never regain it.

After the first round and each subsequent round, the auctioneer will post – possibly partial – information about the bids in the previous round. At a minimum, the auctioneer will publish the minimum required bids for the following round, and usually some measure of the excess demand for each lot. In most of the first SMR auctions, the auctioneer would report all bids and bid amounts.³ Also after each round, the auctioneer informs each bidder of the lots for which it has a provisional winning bid.

³As discussed in more detail below, concerns about signaling and collusive bidding led the FCC and other regulatory agencies to adopt limited-disclosure bidding. See [18] and [15]. Also, [50] discuss bidder collusion in SMR auctions.

In contrast, in many SMR auctions, a bidder can submit jump bids, that is, bids above the minimum required amounts. This can serve to influence the rate at which prices of different lots increase during the auction, as well as the relative prices. A bidder may wish to do so when there are complements and the bidder seeks to determine its ability to win one lot before having to commit to another.

In each round after the first, bidders that were topped on any lot can improve their offer, switch their points to other lots, or reduce eligibility. A bidder failing to maintain its required activity in one round will lose eligibility in the next. Note that, typically, a bidder does not have to increase its offer on a provisional winning bid to have that bid count toward activity.⁴ As prices rise,⁵ bidders will find the cost of keeping eligibility increases, and so will have incentive to reduce activity, and eligibility. Also, if the auction starts with less than a 100% activity requirement, the activity requirement will tend to increase toward 100%. So, as the auction progresses, bidders will tend to drop points and demand. The auction ends when there is no longer excess demand on any lot.

There are many variants of the SMR auction format. Perhaps the most fundamental variant is the simultaneous clock auction, in which bidders do not name prices, and only indicate which lots they want at the announced clock prices. As there is often little incentive for bidders ever to bid much or anything above minimum required bids, the simultaneous clock auction⁶ is really not much different than the SMR. The inability of bidders to submit jump bids, that is, increase offers above the minimum required bids, has two effects. One is that there is less scope for signaling in the clock auction than in an SMR auction. The other is that relative prices are controlled by the auctioneer. The lack of signaling can produce more competitive outcomes. The fact that bidders have less influence on relative prices in a clock auction means that decisions may be sequential – the activity rules require decisions on large lots to occur before decisions on smaller ones, making the auction sequential, which can result in inefficiencies and price anomalies.

SMR auctions vary in many other respects. As noted above, the activity re-

⁴The Indian 3G auction in 2010 was an exception.

 $^{^5\}mathrm{This}$ section assumes a forward auction. In a reverse auction, prices start high and then fall.

⁶See [49] for a description of a simultaneous descending clock (reverse) auction.

quirements and information reported to bidders are auction design features. The effect that these design features can have on the auction outcome is explained below. SMR and clock auctions can be run for multiple units of a single product, single units of multiple products, and multiple units of multiple products. Simultaneous auctions have also been used for procurement.

9.3 SMR Auction Properties

The SMR auction is quite a bit more complicated than a standard English auction or a sealed-bid auction. The question arises as to what value is there in using an SMR or a clock auction format. The SMR auction can, under circumstances described in what follows, effectively allow bidders to arbitrage price differentials. This ensures that the outcome is efficient, unlike the case in sequential auctions.

9.3.1 Straightforward Bidding in SMR Auctions

As has been discussed above, sequential auctions, or simultaneous auctions for substitutes or complements, can result in difficult decisions for bidders and in inefficiencies. The SMR auction gives bidders the ability to arbitrage price differences between substitute products. In other words, a bidder can shift activity from lots for which the value-price differential is low to one for which it is high. Thus, a bidder can bid *straightforwardly*. Note that this requires that bid increments be small, and the activity rule provides flexibility to make them so. For example, suppose lots A and B are very close substitutes for a bidder, but that lot A has more activity points than B. If the prices are about the same, a bidder may prefer A to B.

If price of A is a fixed amount, Δ , more than B, the bidder would prefer B. If the prices start the same, this bidder may start bidding for lot A. However, if the price of A rises faster than that of B, the bidder would want to and can switch to B. Now if, later in the auction, the price of B rises to catch up with, and perhaps pass, A, this bidder would want to switch back. But an activity requirement of 100% (or more than the ratio of the points for B to the points for A) would not permit this.
When such arbitrage can occur during the SMR auction bidding rounds, then, on the margin, a bidder will be indifferent between two lots only when the price-value differential is the same across the two. As all bidders can shift, the auction should end where the marginal bidder is just indifferent. This is a standard criterion for efficiency in competitive markets. More specifically, assuming bidders always bid on the set of blocks that maximize surplus in each round, the outcome of an SMR or clock auction should be both efficient and a *competitive equilibrium*. In other words, at the final prices, excess demand is zero for all lots, which is the criterion for a set of prices to be a competitive equilibrium.

Summarizing, this logic implies the following result [54] (p. 270), which applies to a simplified SMR auction format in which bidders continue to bid on all they want until they have decided to drop eligibility in each round, and in which activity rules allow switching consistent with revealed preferences.

Theorem 3 Straightforward bidding is a feasible strategy for bidder j for all initial prices, all fixed increments, and all feasible price paths iff all goods are substitutes for bidder j.

This theorem does not apply when, as discussed above, the auctioneer does not assign close substitutes the same number of activity points, thus limiting bidders' opportunities for arbitrage. Note, too, that if goods can be complements, and there are two packages, and a bidder gets topped on some, but not all, lots in a package, then it may lack the eligibility to switch to what is a higher-surplus package at the current round's prices. This can happen when, for instance, prices on the lots in one package increase, reducing the surplus that bidders can derive from it, while prices in another package do not. The fact that bidding converges to a competitive equilibrium, which is also an optimum, follows from the fact that straightforward bidding means all bidders are able to arbitrage price differentials. So, if one bidder valued one unit of one product more than a rival who won it, it would always have the opportunity to outbid that rival, and so the auction could not end with a suboptimal allocation, and the prices could never be other than those corresponding to a competitive equilibrium. Summarizing:

Theorem 4 ([54]). Assume that all goods are substitutes for all bidders and that all bidders bid straightforwardly. Then the auction ends within a finite

number of rounds. The final provisional winning bids and allocation of goods are a competitive equilibrium for an economy in which bidder valuations are within some small $\epsilon > 0$ of the bidders' true valuations for all bidders. The final assignment maximizes total welfare within a single bid increment.

Notice too that an SMR auction must converge in a finite number of rounds, assuming fixed increments. This follows from the fact that there is only room for a finite number of rounds before all prices would exceed all bidders' values. Thus, given any initial prices, any SMR auction converges to some limit. Further, given the arbitrage condition with substitutes, this limit must be a competitive equilibrium, assuming all lots are substitutes and bidders bid straightforwardly. Thus, a competitive equilibrium always exists for the case of substitutes.

9.3.2 Strategic Withholding in SMR Auctions

The above subsection shows how SMR and simultaneous clock auctions are efficient when bidding is straightforward and all the lots are substitutes. This subsection looks at the incentives for straightforward bidding and suggests some measures that can improve those incentives.

To illustrate why an individual bidder may not always want to bid straightforwardly, consider an auction in which two bidders are competing for two lots. The first bidder has a valuation for two lots, and the second bidder only wants one. Other bidders are assumed to have lower valuations. Suppose when prices rise to 50, this two-lot bidder, which values each lot at say 100 or more, sees it can win one lot. If this bidder believes it can win both lots for a price not much higher than 50, then it will want to keep bidding. However, if this bidder may have to bid nearly 100 to win both lots, or may not be able to win the second lot at all, then it will want to drop its demand for the second lot as soon as it sees it can lock down the first lot. Even if this two-lot bidder believes it is just as likely to be able to win the second lot at 50 as at 100, it will have no incentive to continue bidding. In other words, this bidder will wish to withhold demand.

This logic readily extends to situations commonly encountered in SMR spectrum auctions in which two bidders are competing in several geographic areas. If each were to accommodate the other, then the auction could end at a low price; on the other hand, should they keep bidding, the auction could end at a high price for the same allocation that might have occurred had the two each reduced demand. Early SMR auctions allowed bidders to submit arbitrary bid amounts. So, a bidder could submit a very large bid, and add a few dollars or euros to bids on some lots over others. This type of bidding, that is, adding a few "trailing digits," can be interpreted as giving signals.⁷

Regulatory authorities now usually specify a limited set of allowed bid increments. Thus, in some auctions, a bidder can only bid a fixed increment, and in other auctions a bidder can choose among a fixed set of increments (e.g., $10,000, 20,000, \ldots, 50,000$ above the provisional winning bid. These nondiscretionary increments limit the possibility of signaling. Of course, overt collusive agreements among bidders are never allowed. So, in the end, signals are never more than cheap talk. But the cheap talk can, at times, affect the outcome.

Moreover, by not reporting the exact demand after each round, the auctioneer can further reduce bidder incentives to withhold demand. In the first example, if the two-lot bidder does not know that there is exactly one lot of excess demand, it cannot know that by reducing demand it can end the auction. Thus, it is now common for auctioneers to limit the information reported between rounds about excess demand.

These limited-disclosure rules can come at a cost. Recall that in a commonvalue auction, aggregate information about rivals' demand can reduce uncertainty about a bidder's own value forecasts. This reduction in uncertainty can result in higher auction revenues. Thus, an auctioneer risks reducing auction revenues when it reports no information about excess demand to any bidders between rounds. At times, auctioneers will report approximate excess demand. For example, rather than indicating that there is only one block of excess demand, the auctioneer can indicate that there are fewer than five blocks of excess demand. Such information can alleviate the winner's curse in common-value auctions. Indeed, if the information reported is sufficient,⁸ in the sense that any additional information would not help bidders improve their value forecasts, then full disclosure will have no additional benefits. In

⁷See [18] and [15] for a discussion.

 $^{^{8}[24]}$ is an early reference for the concept of sufficient statistics.

such cases it may be possible to provide adequate information to bidders to take full account of the winner's curse without having to risk increasing incentives for strategic withholding.

Along these lines is the decision that the FCC and other regulators have taken to limit bidder discretion in setting bid increments. In most of the early SMR auctions, bidders were allowed to submit bids in any amount, and the last few digits of bids may have been used in efforts at covert signaling.⁹ In one case, signaling resulted in a post-auction conviction of a bidder who attempted to collude during an auction with another bidder. Most recent auctions allow a limited number of bid increments. In some cases, the auctioneer allows only yes or no bids for a fixed increment, which turns the auction into a simultaneous clock auction.

One other approach has been used to offset bidder incentives to strategically withhold demand. In the above example, in which there are two lots available in the auction, the two-lot bidder will have less incentive to withhold demand if the auctioneer can decide to only sell one lot. For example, the auctioneer can decide, in advance, to sell only one lot to any single bidder unless that bidder is willing to pay at least 75 or at least 50% more than a rival is willing to pay for the second lot. In this case, the large bidder can still withhold demand, and may still have an incentive to do so, but its incentive is reduced.¹⁰ This can be effective when the auctioneer does not have to sell its entire inventory, and can carry it over to another date. This approach has been used in energy procurement auctions.¹¹ When the auctioneer does not have the ability to carry over the inventory to another date, or other outlets, then this type of volume reduction will not be effective. In some cases, there may not even be the option of just keeping the unsold inventory.

⁹See [18]. Also, see [23] for a different analysis of the impact of information disclosure provisions.

 $^{^{10}\}mathrm{See}$ [51] for a more formal analysis for the case in which the auctioneer does not have to sell both, or all, lots.

 $^{^{11}}$ See [49].

9.4 SMR Auction Experience

The first few auctions were intended to assign rights to narrowband and broadband personal communications services (PCS) licenses. Narrowband PCS was intended for one-way and two-way paging and messaging services. The broadband PCS was for first digital (or 2G) mobile voice and data service. One part of the auction process was to determine the geographic area of coverage and the bandwidth associated with each license. In the first narrowband PCS auction, the FCC divided up the narrowband spectrum into ten national licenses of different bandwidths. In the second, it divided the spectrum into six bands in each of five geographic regions. The broadband PCS was initially scheduled to occur in a sequence of three auctions. The first, FCC Auction 4, included two 30 MHz licenses in each of 51 major trading areas (MTAs). The second, for designated entities (small businesses, minority- and women-owned businesses, and rural carriers), was for another 30 MHz in each of 493 basic trading areas (BTAs), which were finer than the MTAs. The third auction was for three 10 MHz blocks in each of the 493 BTAs.

The very first SMR auction – FCC Auction 1, Figure 9.1, for data communications services – appeared quite successful in that it apparently facilitated arbitrage. The frequencies were divided into three types – five twoway licenses with 50 kHz uplink and downlinks, three two-way licenses with asymmetric 50 kHz downlink channels and 25 kHz uplink channels, and two one-way 50 kHz licenses. Each license with a particular bandwidth sold for essentially the same price as every other license of that type.

The second SMR auction (Figure 9.2) – (which was actually the 3^{rd} FCC auction¹² – included six licenses in each of five regions. Again, the outcome strongly suggests that the SMR auction is an effective mechanism for assigning spectrum licenses. Note that in this case the regional blocks were likely complements. The activity rules apparently provided bidders adequate flexibility to solve their exposure problems.¹³

 $^{^{12}\}mathrm{The}\ \mathrm{FCC's}\ \mathrm{second}\ \mathrm{spectrum}\ \mathrm{auction}\ \mathrm{was}\ \mathrm{not}\ \mathrm{an}\ \mathrm{SMR}\ \mathrm{auction}$

¹³A bid withdrawal in the South region resulted in one of the blocks selling for a significant discount with respect to the others. However, if one adds back the withdrawal penalties that needed to be paid, the total price for this block was very similar to that for the others.

Figure 9.1:

FCC Auction #1 – National Narrowband PCS

Name	Туре	Final Bid	Minority Credit	Winner	Round
N-1	50-50	80,000,000	0	Pagenet	37
N-2	50-50	80,000,000	0	Pagenet	37
N-3	50-50	80,000,000	0	McCaw	33
N-4	50-50	80,000,000	0	McCaw	33
N-5	50-50	80,000,000	25%	MTel	37
N-6	50-12.5	47,001,001	0	AirTouch	24
N-7	50-12.5	47,505,673	0	BellSouth	25
N-8	50-12.5	47,500,000	25%	MTel	24
N-10	50-0	37,000,000	0	Pagenet	45
N-11	50-0	38,000,000	25%	Pagemart	46
Total		671,006,674			

Figure 9.2:

FCC Auction #3 – Regional Narrowband PCS

Band/Region	Northeast	South	Midwest	Central	Western	US Total
50 x 50 kHz	\$17,500,000	\$18,400,000	\$16,810,000	\$17,340,000	\$22,549,020	\$92,599,020
$50 \ge 50 \text{ kHz}^*$	\$14,850,000	\$18,780,000	\$17,360,401	\$17,136,000	\$22,800,000	\$90,926,401
50 x 12.5 kHz	\$9,471,082	\$11,800,007	\$9,291,000	\$8,250,000	\$14,857,003	\$53,669,092
50 x 12.5 kHz	\$8,949,543	\$11,543,007	\$10,057,004	\$8,791,001	\$14,281,111	\$53,621,666
50 x 12.5 kHz	\$8,675,000	\$8,000,013	\$9,500,000	\$8,262,000	\$14,281,001	\$48,718,014
50 x 12.5	\$10,251,000	\$11,262,003	\$10,251,001	\$10,488,000	\$10,920,600	\$53,172,604

Figure 9.3:

UK 3G Auction Results

License	High Bidder	Final Price £m
А	TIW(Hutchison)	4,348.70
В	Vodafone	5,964.00
с	BT3G	4,030.10
D	One2One	4,003.60
E	Orange	4,095.00

The SMR auction format has been used by regulatory authorities for a great many auctions for spectrum and other products and assets across the world since the mid-1990s. For the first wave of European spectrum auctions, the SMR format was used in Austria, Germany, Italy, the Netherlands, Switzerland, and the UK. In virtually every case, similar licenses sold for similar prices, and larger licenses sold for more than smaller ones. All these auctions included 120 MHz of paired UMTS spectrum, which could be divided into 4, 5, 6, or 12 lots.

The UK and Italy decided to divide the spectrum into three lots of 20 MHz and two lots of 30 MHz. One of the larger lots was reserved for an entrant.¹⁴

The UK was the first European country to conduct a 3G auction. That auction attracted the four incumbents (BT, Vodafone, Orange, and One2One) and nine potential entrants. Each firm could bid for one license and had to remain active to remain in the auction. The four incumbents each won a license (and Vodafone won the larger, 30 MHz license). H3G won the license reserved for the entrant. The auction results were as in Figure 9.3.

 $^{^{14}}$ The 120 MHz was actually split into separate uplink and downlink channels, and so a 20 MHz license typically included a 10 MHz uplink and a 10 MHz downlink channel, and is sometimes written as 2×10 MHz, or 10 MHz for short. In addition, the UK and Italian auctions included some other, less valuable spectrum for what are called time division duplex (TDD) applications.

German 3G Auction

Blocks	Bidder Name	DM	USD
1	Viag Interkom:	DM8.3104 billion	\$3.86 billion
2	Mobilcom Multimedia:	DM8.17 billion	\$3.79 billion
3	Mannesmann Mobilfunk:	DM8.33 billion	\$3.87 billion
4	Group 3G:	DM8.3046 billion	\$3.85 billion
5	Mobilcom Multimedia:	DM8.2 billion	\$3.81 billion
6	Viag Interkom:	DM8.2066 billion	\$3.81 billion
7	T-Mobil:	DM8.3043 billion	\$3.85 billion
8	E-Plus Hutchison:	DM8.2743 billion	\$3.84 billion
9	T-Mobil:	DM8.2779 billion	\$3.84 billion
10	E-Plus Hutchison:	DM8.1439 billion	\$3.78 billion
11	Mannesmann Mobilfunk:	DM8.1438 billion	\$3.78 billion
12	Group 3G:	DM8.1414 billion	\$3.78 billion

Figure 9.4:

The German 3G auction was designed differently. The same 120 MHz that was also available in the UK was divided into 12 blocks of 10 MHz. Each bidder was allowed to bid for two or three blocks. A bidder that was the high bidder on one block at the end would not win, nor would it owe anything, and the unsold block would be immediately re-auctioned. The regulator did not report all bids, but only the identity of standing high bids at the end of each round.

There were seven bidders that applied to participate in the German auction: the four incumbents – T-Mobile, Mannesmann (Vodafone), Viag (O2), and E-Plus – as well as three potential entrants. The auction did not reserve any blocks for entrants, and indeed was criticized beforehand for allowing the incumbents to foreclose entrants. The auction thus allowed for four, five, or six winners.

One of the entrants, Debitel, dropped out when the price reached DM¹⁵ 4,897 million per block in round 121. This was not common knowledge for all the bidders, as only high bids, and not all bids, were reported by the regulator after each round. However, bidders could guess after a few rounds that Debitel had dropped out. Once it had done so, the other bidders could settle right away for 20 MHz, or keep bidding, hoping to win 30 MHz and face less competition in the market after the auction.

 $^{^{15}\}mathrm{The}$ Deutschemark was worth about 0.5 euros. It was soon phased out.

Mannesmann and T-Mobile elected to keep going for another 53 rounds, until prices reached DM 8,330 million. At that time T-Mobile found it more prudent to drop out, and risk ending with two blocks and a six-player market, than keep going. Had Mannesmann won three blocks, there would have been one unsold block, which T-Mobile could have tried to acquire in a second auction.¹⁶

The results of the German auction are in Table 9.4. Each of the six winning bidders won two blocks.

This was not the real end of the process. Neither of the entrants that *won* spectrum, Mobilcomm (backed by France Telecom) and Group3G (back by Telefonica) ever elected to roll out service. Both abandoned their licenses, and for no refund or other considerations. Germany was left with a four-player market.¹⁷

The Italian auction (Figure 9.4) used the UK design; it apparently also achieved a reasonably efficient outcome despite the very limited competition – there were six bidders for five licenses.

The Austrian auction copied the German design, but there were six, rather than seven, bidders to start. That auction settled in seven rounds.

The SMR and simultaneous auction format has now been used in dozens of auctions for spectrum – almost all the spectrum auctions in both North and South America, Australia, and parts of Asia, including India, Singapore, and Taiwan. It has also been used for a number of energy auctions; perhaps the first successful application was the simultaneous descending clock auction

¹⁶[41] account is a bit different. He indicated that T-Mobile and Mannesmann should have either dropped out two blocks earlier or kept going until one of the entrants dropped out. This would have been a sound pre-auction recommendation if T-Mobile and Mannesmann had both believed that the entrants were either weak or quite strong. As it was, each bidder faced a great deal of uncertainty about rivals' valuations, and as the two strongest incumbents, neither T-Mobile nor Mannesmann wanted to be the first to cede the third block, and possibly lose market share. Only when prices rose so high that a re-auction could be expected to result in lower prices for the unsold extra block did T-Mobile find it advantageous to reduce eligibility.

 $^{^{17}}$ These four blocks were re-auctioned 10 years later for a total of EUR 350 million, or less than 2.5% of the earlier auction price. Moreover, Telefonica spent approximately EUR 1 billion on network buildout before deciding to walk away.



Italian 3G Auction



(SDCA) used for energy procurement in New Jersey, 18 and then Italy, Ohio, Illinois, and elsewhere, and for various other energy products in Austria, France, and Germany. 19

9.5 SMR Auction Configuration

Experience with the SMR auction format has not been uniformly successful. What follows provides some examples of outcomes that appear inefficient – in some cases larger and more valuable lots sold for less than smaller lots. This section explains why at times different bidders pay different prices for identical aggregate amounts, and why bidders winning more can even pay less. In general the auction prices and allocations appear to leave some bidders overpaying relative to others and suggest that misallocations or inefficient allocations occurred.

Here are several examples of how auction configuration (and in particular activity rules are implemented), and in combination with bidding strategy, can affect the outcome of an auction.

¹⁸See [49].

 $^{^{19}}$ See [6].

	Reliant South	TXU South	Δ	$\%\Delta$
1 year base load	\$7.32	\$10.59	\$3.27	45%
2 year base load	\$6.59	\$10.33	\$3.74	57%
1 year base cyclic	\$0.75	\$1.79	\$0.95	127%

Table 9.1: Texas capacity auctions: Fall 2001

9.5.1 The Texas Capacity Auctions

Perhaps the most glaring example of what can go wrong is seen in a few auctions for generation capacity in Texas in 2001 (Table 9.1 and 2002 (Table 9.2).²⁰ The first capacity auction included one-year and two-year contracts.

These auctions were for 25 MW blocks of capacity for specific durations. The capacity was owned by three companies: CPL, Reliant, and TXU. These generation companies were under a regulatory mandate to sell off a fraction of their capacity, which they did in a series of quarterly auctions starting in 2001. The tables here are for the first two auctions.

In the Fall 2001 auction, the Reliant and TXU capacity was for *identical* blocks from the same physical base load generators jointly owned by Reliant and TXU. These blocks were sold in an SMR auction. In that auction, Reliant and TXU separately managed the starting prices and bid increments, and assigned separate eligibility points for each "auction." Reliant started with a lower price, but TXU set lower increments. The result was that the Reliant part of the auction cleared first. When the TXU price subsequently passed the Reliant price, bidders that had stopped bidding on the Reliant product were stuck. There was no provision in the auction rules to allow bidders to switch back and forth.

The second capacity auction was for monthly contracts. A similar pattern arose in that subsequent Spring 2002 auction. Again, the CPL and Reliant products were identical. The bids were for capacity, but fuel costs were added in separately. Bidders failed to adapt in the second Texas capacity auction. Eventually, the Texas Public Utility Commission mandated that a switching

²⁰See http://www.puc.texas.gov/agency/rulesnlaws/subrules/electric/25. 381/24492adt.pdf

	Reliant South	CPL South	Δ	$\%\Delta$
June base load - capacity only	\$13.13	\$6.85		41.3%
June base load - capacity + fuel	\$22.57	\$18.76	\$3.81	
July base load - capacity only	\$6.76	\$10.60		45.3%
July base load - capacity + fuel	\$22.57	\$18.76	\$3.06	
August base load - capacity only	\$16.76	\$9.85		55.6%
August base load - capacity + fuel	\$26.86	\$22.79	\$4.07	

Table 9.2: Texas capacity auctions: Spring 2002

rule be introduced.²¹

9.5.2 The Spanish 4G Auction (2011)

The 2011 Spanish 4G auction provides another striking example of what problems can arise with the activity requirements in SMR auctions which include both regional and national licenses.

This auction (Figure 9.6) included nine bands within the 2.6 GHz band. Seven were for national licenses (four each of 2×10 MHz, and three each of 2×5 MHz), and the other two were divided into 19 regions (one license of 2×10 MHz and one of 2×5 MHz in each region). The three largest bidders each won 2×20 MHz. However, there were many ways for this outcome to emerge. What happened was that Vodafone won a combination of national 2×5 MHz and regional 2×5 MHz licenses, whereas the two other incumbents, Orange and Telefonica, each won national licenses. The main decision facing these operators in this auction was when to switch from the national to the regional licenses. The activity requirements made it almost impossible to switch back. Until one switched to regionals, there would also be little competition to cause prices to increase. So the decision to switch had to based on an estimate of how hard the regional operators would compete to get the 2×5 MHz blocks in each region. The outcome suggests that Vodafone, which ended up switching first, guessed incorrectly, as it paid almost 30%more than its rivals. (See Table 9.3)

²¹See http://www.puc.texas.gov/agency/rulesnlaws/subrules/electric/25. 381/24492adt.pdf .

		Price	e of 2 x 10	Prie	ce of 2 x 5	
Coverage	Winner	MHz		MH	lz	Totals
National	Orange	€	22,214,958			€ 45,096,364
National	Orange	€	22,881,406			
National	Telefonica	€	21,791,816			€ 44,438,132
National	Telefonica	€	22,646,316			
National	Vodafone			€	9,811,904	
National	Vodafone			€	9,789,654	
National	Vodafone			€	10,150,457	€ 59,095,246
Regional	Vodafone			€	29,343,231	

Figure 9.6: Spanish 4G auction (2011)

What is also striking in the Spanish auction is that in each of the three most expensive regions, Catalunya (including Barcelona), Madrid, and Pais Vasco, the price of the 2×5 MHz license was more than the price of the 2×10 MHz license.

Region	Population	Price of a 2x10 MHz license	Price of a 2x5 MHz license				
Cataluña	7,512,381	€4,239,078.48	€5,888,062.90				
Madrid	$6,\!545,\!684$	€3,818,047.06	5,043,280.86				
Pais Vasco	$2,\!178,\!338$	€2,417,683.79	$3,\!186,\!291.36$				

Table 9.3: Spanish 4G auction - 5 MHz vs 10 MHz prices in 3 large regions

This was also a result of the limitations imposed by a combination of the activity rules and spectrum caps. The activity rules did not allow a bidder to easily arbitrage between the prices of the national and regional licenses. And there were spectrum caps that did not permit a bidder to bid for 2×10 MHz when 2×5 MHz was more expensive.

	MPOPs	POPs	PWB total	Ave \$/MPOP
Cricket + Denali	2,201,624,760	176,183,964	\$ 985,082,750	\$0.447
Spectrumco	5,287,189,470	$267,\!387,\!437$	2,337,609,000	0.451
Cingular	2,436,458,880	198,768,198	1,334,610,000	0.558
T-Mobie	6,638,718,070	474,718,308	\$ 4,182,312,000	0.630
Cellco(VZW)	3,840,952,220	192,047,611	\$ 2,808,599,000	0.731
MetroPCS	1,445,444,020	$144,\!544,\!402$	1,391,410,000	0.963

Table 9.4: FCC Auction 66: Final prices paid by 6 largest winners

9.5.3 The Mexican PCS Auction

Another quite striking example of how activity rules can limit arbitrage across difference size licenses, in this case of 10 MHz and 30 MHz licenses in the Mexican 2G auction (Figure 9.7). The activity rule specified that the 30 MHz licenses counted for five times as many points as the 10 MHz licenses. This made switching difficult in one more way: a bidder dropping from 30 MHz to 10 MHz might not be able to switch back. The next effect was that in one region, 7, the 30 MHz licenses sold for more than the 10 MHz licenses.

Activity points that had permitted bidders to more easily arbitrage the price differentials between the 10 MHz and 30 MHz licenses would have prevented this type of anomaly from occurring. It is not clear that there was a benefit in assigning more points to the 30 MHz blocks than the 10 MHz ones in any region. The goals of the activity rules include facilitating efficient arbitrage and also compelling bidders to mark increasingly serious offers. Assigning a larger number of points to the larger regions probably accomplishes this goal.

9.5.4 US FCC Auction 66 – Advanced Wireless Services

The US AWS auction, Auction 66, was another auction in which activity rules limited arbitrage (Table 9.4).

In that auction, the FCC allocated six bands of spectrum partitioned geographically in three ways: one 20 MHz block was allocated in each of 734

Figure 9.7: Mexico 2G Auction Results

		Block (MHz)				
Region	Population	A(30)	B(30)	D(10)	E(10)	
1 - Baja	26,216,350	\$161,544,000	\$162,523,000	\$152,400,000	\$137,870,000	
2 - Sonora	4,375,593	\$29,560,000	\$31,642,000	\$15,000,000	\$14,870,000	
3 Chihuau	4,978,464	\$101,656,000	-	\$70,100,000	\$70,000,000	
4 - Monte	7,493,285	\$616,877,000	\$588,488,000	\$363,300,000	\$353,252,000	
5 - Campe	8,256,749	\$9,000,000	-	\$6,800,000	\$6,700,000	
6 - Guada	10,934,908	\$289,699,000	\$286,811,000	\$89,650,000	\$81,077,000	
7 - Agau/0	10,339,977	\$70,876,000	\$75,941,000	\$82,800,000	\$82,121,000	
8 - Oaxaca	18,382,181	\$20,397,000	\$20,160,000	\$6,300,000	\$6,400,000	
9 - DF	23,742,926	\$892,500,000	\$941,506,000	\$485,100,000	\$479,160,000	

cellular market areas (CMAs), one 20 MHz and one 10 MHz in each of 176 economic areas (EAs), and one 20 MHz and two 10 MHz licenses in each of 12 regional economic area groups (REAGs). The average prices were highest for the REAG licenses²²; they were over \$0.10 (20%) more than for the CMA and EA licenses. Also, the average price by bidder varied quite a bit. Auction 66 had such a range of licenses that assigning points to prevent activity traps very difficult. To address this issue some recent (combinatorial clock) auctions ("CCA"s), a *revealed preference* activity rule has been imposed. A reveal preference rule requires bids in any one round be consistent with those in previous rounds.²³

While each bidder won a different footprint, some comparisons can be made. T-Mobile won 20 MHz licenses in approximately 80% of the top 100 markets. In contrast, Spectrumco covered essentially all of the top metro areas and paid approximately 25% less than T-Mobile. Cellco and Cingular combined to cover the entire US; Cricket and MetroPCS also covered all the major metropolitan areas.²⁴ Clearly, as in the Spanish 4G auction, the choice of whether to purchase the larger REAGs or the smaller EAs or CMAs affected the average price. In contrast with Spain, the prices in this auction for the large regions were higher.

The outcome of Auction 66, while perhaps difficult to predict, can be explained at least partly by the bidding strategies. The two parties achieving the lower average prices (the Cricket-Denali consortia and Spectrumco) were apparently better able to determine when to switch from the larger REAGs down to the smaller - 0.451 for Spectrumco and 0.630 for T-Mobile. There were two approaches used to make this determination.²⁵ Spectrumco employed jump bids to try to assess the budgets of at least marginal bidders. In many auctions, the total exposure – that is, the sum of the values of all bids, both high bids and bids that have been topped – can peak long before the

²²See http://wireless.fcc.gov/auctions/default.htm?job=auction_ summary&id=66

 $^{^{23}}$ This type of rule is an application of the logic of the axiom of revealed preference (see [66]). The CCA is discussed in more detail below.

 $^{^{24}}$ See also [16] for a discussion of the bidding strategy that recognized the flaw in the activity rule. Bulow et al. did not include Cricket in their report. Salant [65] also provides an example of how a bidding strategy can be adapted to take advantage of rivals' budget limitations and limited arbitrage opportunities in an auction.

²⁵See [16] for a discussion of Spectrumco's strategy. I advised Cricket.

auction closes. This peak can be used to provide an estimate of the budgets of all bidders in the auction.

If one bidder *knows* its rivals' budgets, then it can use this information to achieve a better allocation and lower prices.²⁶ Consider the following example (9.5.4):

Two stuget constrained studers and timee lots					
Bidder/Value for	Lot 1	Lot 2	Lot 3	Budget	
Bidder A	100	100	100	100	
Bidder B	100	100	100	100	

Two budget-constrained bidders and three lots

Table 9.5: Two bidder and three lots

In this example there are two bidders and three lots. Each bidder has values for each lot of 100 (or more) but is quite budget constrained, having a total budget of 100. If one bidder can get its rival to bid 50 or more for one lot, it can win the other two lots. And the more it can get that rival to sink in the one lot, the lower will be the price it has to pay for the other two. The multi-round structure of an SMR auction allows for this type of strategy. It is the type deployed by Spectrumco in Auction 66.

This budget-based strategy has some risks. It can be difficult to forecast budgets. And even the ability to forecast budgets is no guarantee of retaining the flexibility to take advantage of it. Moreover, it is unusual for all bidders to face effective budget constraints.

Another approach in Auction 66 and in the Spanish auction is to try to forecast the conditions under which the small blocks will sell for less than the larger blocks. This is the approach taken by Cricket and Denali. If there are both large bidders seeking larger regional licenses and small bidders seeking only a few licenses, then the large bidders will want to stay on the larger blocks until they get to be more expensive than the sum of the smaller blocks. At some critical premium, the large-block bidders will want to shift to regional blocks, anticipating that they will end up being a better value. When this should occur will depend on the relative amounts of competition

 $^{^{26}}$ See [65] for a more detailed discussion.

for the larger blocks and the smaller blocks. To some extent, this can be observed during an SMR auction.

9.5.5 Arbitrage in SMR Auctions

The above examples indicate that arbitrage opportunities can be limited in SMR auctions. These examples raise two immediate questions. One is whether there is harm – that is, whether there is a loss of revenue for the auctioneer, or bidders are harmed. The other is what measures can be taken to prevent any such harm.

Proposition 2 Consider an SMR auction with at least two bidders and two lots in which activity rules do not allow bidders to switch between substitutes, at least in some stages of the auction. Then, for some preference profiles, this auction will necessarily result in inefficient outcomes even with straightforward bidding.

Proof: Suppose there are at least two lots, or combinations of lots, and at least two bidders, and that one bidder that was topped on one block is no longer eligible to switch to the other. If the auction were to end, then there would be an alternative allocation in which this bidder would get the other block and the bidder with the second block would not. This alternative allocation would result in higher bidder payoffs and revenues whenever the bidder that should not win has a lower value for that second block than the bidder that should win. \Box

Recall Milgrom's result that the SMR auction with straightforward bidding will result in an efficient outcome. This results in certain conditions about activity requirements. What the above proposition means if that the activity rules do not permit arbitrage, then the outcome can be inefficient.

Now, there are many ways in which the activity rules can limit arbitrage possibilities.

1. Switching rules. The Texas capacity auctions did not allow any switching. Once Texas introduced switching rules, the price discrepancies disappeared.

9.5. SMR AUCTION CONFIGURATION

- 2. Abstract blocks. Other auctions, e.g., the Spanish and German 4G auctions and the German 3G auctions, assign separate labels and auction prices to various identical abstract blocks. Clearly, combining offers for all identical abstract blocks (assuming an SMR auction or having a single price for all the blocks in a simultaneous clock auction) can only improve efficiency.
- 3. Starting prices, bid increments and activity parameters In simultaneous clock auctions, relative prices are determined by the auctioneer. The starting prices and bid increments determine how fast the prices of different lots increase. What often happens is that the larger lots tend to clear first. This tends to turn a simultaneous auction into a sequential one, as bidders' decisions are essentially sequential.²⁷

For example, in FCC Auction 66 there were three types of geographic licenses with 20 MHz and covering New York City: a 20 MHz Northeast REAG had a bit less than 48,000,000 bidding units, the 20 MHz New York EA had just over 24,000,000 bidding units, and the 20 MHz New York CMA had approximately 16,000,000 bidding units. An activity requirement of 60% would not permit switching from the New York CMA to the New York EA license, but one of two-thirds would. And this would also allow switching from New York and either the Boston or the Philadelphia EA to the Northeast REAG covering those EAs. Thus, the choice of activity requirements can facilitate or hinder arbitrage.

4. Activity points As the above makes clear, the way activity points are set, in combination with the activity requirement percentage, can affect auction efficiency. Generally, when a number of small blocks combine to make one large block, then the small blocks should have a combined number of activity points equal to that of the larger one. Such a provision allows a bidder to switch from a sent of small blocks to a large block with same bandwidth and geographic coverage. And this facilitates arbitrage. The above suggests that the way points are assigned to the small blocks, in combination with the activity percentage requirements, will affect auction efficiency.

 $^{^{27}}$ The Regulatory Authority managing the auction can try to set increments based on activity, as the FCC does a variable price increment rule, that is, one in which the size of the increment is based on bidding activity. See [18] for a discussion.

5. Size of licenses and blocks When there are both regional and national licenses in the same auction, it will not usually be the case that the sum of the regional prices will equal the price of a national license. Similarly, when there are smaller and larger blocks, such as 10 MHz and 20 MHz licenses, it will not usually be the case that the price of the smaller blocks sum to exactly the price of the larger one. Any activity rule, except for one allowing

package bids or one based purely on revealed preferences, will limit arbitrage possibilities, The latter would require that bidders be consistent in how they switch.²⁸

9.6 Conclusion

This chapter describes perhaps one of the most significant contributions economists have made – the simultaneous multiple auction. This auction format has the potential to offer significant improvements both in auction revenues and for bidders, by facilitating better matches of bidders and lots. However, there are a number of ways in which an SMR or simultaneous clock auction can be misapplied and result in an undesirable outcome. This chapter has provided some description of the main SMR auction pitfalls, and some suggestions for avoiding them. The SMR auction is best suited for auctions of substitutes and some types of complements.

Package bidding options are described in Chapter 10.

164

 $^{^{28}\}mbox{Variations}$ of such rules have been tried in some recent combinatorial clock auctions, which are discussed in the next chapter.

Chapter 10

Combinatorial auctions

Summary

This chapter looks at package bidding, or combinatorial, auctions. Bidders in multi-object auctions will at times be bidding for substitutes and/or complements. This gives rise to significant additional complexity for both bidders and the auctioneer. When bidders are considering substitute licenses or packages, they will want to obtain the best package for the best price. This need not happen, and both bidders and the auctioneer can suffer. When there are complements, a bidder seeking a large package may find itself facing an exposure problem, that is, it may end up winning part of the desired package and pay more than it is worth. This problem can deter bidders from bidding close to value. And when a small bidder is competing against a package bidder, the auction rules can tilt the outcome toward the package bidder or the small bidder.

10.1 Introduction

This chapter looks at package bidding, or *combinatorial*, auctions. Bidders in multi-object auctions will at times be bidding for substitutes and/or complements. This gives rise to significant additional complexity for both bidders

and the auctioneer. As seen in the previous chapters, the SMR and clock auctions are very well suited for achieving efficient outcomes when all bidders view all objects as substitutes. However, these auctions can still leave bidders facing some difficult decisions when there are complements, at least for some lots and some bidders in the auction.

When bidders are considering substitute licenses or packages, they will want to obtain the best package for the best price. This need not happen, and both bidders and the auctioneer can suffer. When there are complements, a bidder seeking a large package may find itself facing an *exposure* problem, that is, it may end up winning part of the desired package and pay more than it is worth. This problem can deter bidders from bidding close to value. And when a small bidder is competing against a package bidder, the auction rules can tilt the outcome toward the package bidder or the small bidder.

To better illustrate the issue, suppose there are two lots, and that there are one large bidder and several small bidders for each of the lots. And suppose these two lots are being sold in an SMR or a simultaneous clock auction.

Let L denote the large bidder's value for the two lots, let s^{j1} denote the highest value among the small bidders for lot j = 1, 2, and let s^{j2} denote the second highest value. Suppose that the large bidder has zero value for one lot.

If $L > s^{11} + s^{21}$, then the large bidder should always win. But this may not happen if, for example, the difference between L and $s^{11} + s^{21}$ is not very large or if the large bidder is very reluctant to risk winning only one lot. This is a simple example of the *exposure* problem. Absent package bids, the large bidder has to take a risk to have any chance at winning.

A more complex case arises with three bidders and three lots (Table 10.1):

10	DIC 10.1. 11	fice bluet,	unce for auction
Bidder/Value for	Lots $1\&2$	Lots $2\&3$	Lots 1&3
Bidder A	1	0	0
Bidder B	0	1	0
Bidder C	0	0	1

Table 10.1: Three-bidder, three-lot auction

Bidder A only wants the package of lots 1 and 2, B only wants the package

of lots 2 and 3, and C only wants the combination of lots 3 and 1.

In the first, two-unit example, if the auctioneer were to call out prices for lots 1 and 2, p^1 and p^2 , for which $L \ge p^1 + p^2$ and for which $p^j > s^{j1}$ for j = 1, 2, then the large bidder would be willing to purchase both lots, and the small bidders would not be willing to purchase any lots. In this case, these prices would be a *competitive equilibrium* in that, when all bidders act as price takers, demand exactly equals supply.

In the second example, with three lots and three bidders, there is no competitive equilibrium. To see this suppose the contrary. Let p^j , j = 1, 2, 3, be such an equilibrium. Then, the following conditions must hold:

$$p^j \ge 0 \tag{10.1}$$

for all j = 1, 2, 3. This says that prices cannot be negative. Also,

$$p^j + p^k \le 1 \tag{10.2}$$

for $j \neq k$, j, k = 1, 2, 3. This says that no bidder will pay more than value. But for there not to be any excess demand, (10.2) has to equal 1 for all bidders But this cannot hold for any pair of bidders.

Package bidding allows bidders to place all-or-nothing bids on a set of lots. Thus, in the above example, bidders can place a single all-or-nothing bid on a pair of lots. Bidder A can place a bid of 1 for the package of lots 1 and 2, B can bid 1 for 2 and 3, and C can bid 1 for 3 and 1. This is an equilibrium, with one bidder randomly winning its bid.

However, package bidding is not a perfect solution. First, if there are small bidders whose values add up to more than that of a large bidder, then these small bidders might need to find a way to coordinate their bids. This is the *threshold* problem. To be more specific, consider the case in which $L < s^{11} + s^{21}$, and there are separate bids for lot 1, lot 2, and the package of both lots. If the standalone "winning" bids for lots 1 and 2 together add up to less than L, then the small bidders would each need to raise their "winning" bids to top L. However, each small bidder would want the other to raise its bid more, and absorb a larger share of the gap. This free-riding problem might prevent the small bidders from winning even when they have the higher valuations.

Package bidding also is inherently more complex than restricting auctions to those in which there is a price for each lot, or product, in the auction. If there are n lots available, there will be 2^{n-1} available packages. Other than for very small numbers of lots, or when the types of complementarities are well understood and limited, package bidding can quickly become computationally infeasible. The process for expressing preferences (how can a bidder communicate values for all such packages?) and then finding the winning set of offers can be overwhelming.

This does not mean that package bidding is never helpful. What follows provides some specific types of package auctions, some of which have been used in practice.

10.2 Hierarchical Package Bidding

Perhaps one of the simplest forms of package bidding is one in which bids are allowed on individual lots and on a few predefined packages. This has been called *hierarchical package bidding* (HPB).¹ The US Federal Communications Commission (FCC) conducted one HPB auction, Auction 73,² that raised over \$18.9 billion. In Auction 73, the C block band was divided into 12 regional economic area group (REAG) licenses. Bidders could bid on individual REAG licenses or the package of all 12 REAGs.

The HPB auction was otherwise virtually the same as an SMR auction. Licenses were assigned points. Bidders had to establish initial eligibility by placing upfront deposits for their maximum desired number of points prior to the auction. The HPB auction imposed similar activity requirements: high bids plus new bids had to exceed a given fraction of the eligibility, or the bidder lost eligibility in proportion to the gap between actual activity and the activity target. Prices increased from one round to the next for all licenses for which there was excess demand. And the activity requirement

 $^{^{1}}$ See [29].

²See http://wireless.fcc.gov/auctions/.

was increased in stages as prices rose and aggregate eligibility dropped during the course of the auction. The closing rule was also the same: the auction would only close on all licenses at the same time and only when there was no excess demand for any license or package.

There were a few differences. First, a bidder could place a bid on a license and on a package which contained that license. However, such a bidder would not get credited with double the activity. Secondly, if there was more bidding activity on the packages than on the component licenses, it was possible that the provisional winning bid amounts after a round for a package exceeded the sum of the provisional winning bids for the individual licenses in the package. In this case, some rule was needed to determine minimum required bids for the component licenses, especially where there was no excess demand for some of those components. And bidders bidding for packages were not permitted to bid the component prices if those were lower. The FCC developed a formula for allocating the gap in determining minimum required bids for the component licenses.³ Finally, a bidder who had stopped bidding on one of the component licenses might find that, after several rounds of bidding in which other components in the package prices increased, their old previously losing bid could turn into part of a provisional winning bid. That bidder might have lost eligibility, or used that eligibility for other licenses. The SMR rules were modified to allow that bidder to resume bidding on that component, but not otherwise increase eligibility.

[29] conducted laboratory experiments, with Caltech students acting as bidders, comparing the SMR, HPB, and unconstrained package auctions. In their experiments they found that the HPB allowed bidders in some cases to solve the exposure problem, and therefore achieved greater efficiency and revenues than the SMR auction. Note that the efficiency measure was based on the difference between the actual sum of the values of the final assignment and the maximum possible sum. They also found that unconstrained package bidding offered no additional benefits.

These experiments *assumed* a particular structure of complementarities. They did not allow for arbitrary or random complementarities. Nor did they identify when the threshold problem might become more significant than the exposure problem. What these experiments do suggest is that where there

³See [29] for details.

is a well-defined exposure problem, i.e., different bidders have similar views of the complementarities, then the HPB seems to be an effective and efficient mechanism. In the actual auction the package bidder won the national license.

[29]'s experiments were conducted on behalf of the FCC in advance of FCC Auction 73. The FCC had previously conducted experiments on a different form of package bidding, called the *simultaneous ascending auction with package bidding* (SAAPB).⁴ The results of those experiments showed improvements over the SMR auction similar to those with the HPB. However, the SAAPB did take significantly longer to run than the SMR. Perhaps because SMR auctions can already be quite long (commonly several weeks, but sometimes over four months), the FCC elected to use the HPB over the SAAPB. No tests were run to compare how many rounds an SAAPB takes on average compared to the HPB.

10.3 VCG Auction Properties

The multi-object version of the VCG auction introduced in Chapter 4 is one type of package auction. As was explained in Chapter 4, the VCG has the property that bidders will always have incentives to report true values. Thus, the outcome of a VCG auction should be efficient. However, as is explained in this section, the VCG auction has a number of other features that could block its adoption.⁵

This section also provides a few examples illustrating these undesirable properties. Note that bids in a VCG auction need to convey to the auctioneer the bidders' valuations for all possible combinations of lots. Thus, an individual bid is a set of reported values for each package, or combination of lots, that is available in the auction. In general, this list can be quite large – with n lots, the number of packages is $2^n - 1$. So, the VCG auction may be somewhat impractical if the number of lots is large. But, as the following examples illustrate, the VCG has other drawbacks even where there are few lots for sale (Table 10.2).

 $^{^{4}}See [59].$

⁵Much of this section is from [20] and [19].

10.3.1 Revenues in VCG auctions

One difficulty with VCG auctions is that they may fail to generate much revenue. To see this, recall that bids in a VCG auction are all-or-nothing offers for packages. The winning bids are those mutually consistent bids that maximize the sum of the offers, and not necessarily the auction revenues. The reason is that the price a bidder pays is the difference between the sum of the rivals' reported offers for what they would have won without this bidder, and their reported offers for what they actually do win. The following examples illustrate how revenues are determined, and in some cases, can be zero.

First, consider the following example (Table 10.2) in which there are two objects and three bidders. In this example, it is assumed that bidders can place bids for one and/or both lots:

Table 10.2: Three-bidder, two-lot VCG auction

Bidder/bid for	1 lot	2 lots
Bidder B1	100	100
Bidder B2	100	100
Bidder B3	0	50

In this example, bidders B1 and B2 each have only one of the two lots available. They are each willing to pay 100 for it. Bidder B3 only wants both lots or none, and will pay up to 50. As both B1 and B2 individually block B3 – that is, B3 loses even if either B1 or B2 does not bid – then the VCG prices for B1 and B2 are each zero. So, low revenue can be a feature of VCG auctions.

Simple variations of this example indicate that the VCG has other not so appealing features. First, consider the case in which values are now lower for B1 and B2 (Table 10.3):

Now, B3 wins both lots. So, if B1 and B2 were really weaker than B3, but both overbid (or colluded to do so), then they could win both lots for much less than their value. [20] also explained how the outcome is sensitive to mergers and demergers. If, for example, bidders B1 and B2 were to merge and to be treated as one bidder which wants both lots, then the merged entity would win both lots at a price of 50 rather than 0. So, a merger works

Bidder/bid for	1 lot	2 lots
Bidder B1	20	20
Bidder B2	20	20
Bidder B3	0	50

Table 10.3: Three-bidder, two-lot VCG auction

to increase the amount that a bidder may have to pay, and a demerger can reduce it. Indeed, rather than splitting into two entities, a combined B1-B2 bidder can enter a second shill bidder. By doing so, B1-B2 can reduce the price it has to pay. This also raises the question whether the auctioneer would want to stop the shill from bidding. In most auctions, more bidders result in higher prices paid. But this need not be the case with VCG auctions. Competition can lower price.

10.3.2 Fairness in VCG auctions

The VCG auction may result in efficient but unfair outcomes. Arguably, one bidder might be troubled if it wins less than a rival bidder but pays more. The following examples show that this can happen in theory and has happened in practice.

Consider the case of two bidders and three lots (Table 10.4):

10	JOIC 10.	I. I WO	bluuci,		٧U
Bidder/bid for	1 lot	2 lots	3 lots	_	
Bidder B1	100	200	280	_	
Bidder B2	100	120	130		

Table 10.4: Two-bidder, three-lot VCG auction

The Vickrey price a bidder pays will be the difference between the value its rivals get at the final allocation and what the rivals would have won had that bidder never submitted a bid. In the above table, there are three lots for sale, and two bidders. Each bidder submits bids for one, two, or three lots. In this example, bidder B1 wins two lots and B2 wins one. The value of the two lots to B1 is 200. B1 also indicated that all three lots together were worth 280. B2 wins one lot worth 100, but would have paid 120 for two lots and 130 for all three.

Thus, B1's VCG price is 30, as B2 only offered an additional 30 for the two lots that B1 wins. And B2's VCG price is 80. This is clearly a bit odd, and B2 may view it as unfair. VCG auctions will not generally have the property that a bidder winning more than a rival pays less. Indeed, the reverse can occur. Notice too that that the outcome is efficient in that it maximizes the sum of the bidder values across all feasible allocations.

Figure 10.1 shows that this unfairness is not just a theoretical possibility. It shows the outcome of a Swiss spectrum auction.

The first four categories (800 MHz, 900 MHz, 1800 MHz, and 2100 MHz) were the most valuable, and the 800 MHz and 900 MHz were more valuable than the 1800 MHz or the 2100 MHz. Each block was 10 MHz, and there were six in the 800 MHz, seven in the 900 MHz, the equivalent of fifteen in the 1800 MHz, and twelve in the 2100 MHz category.⁶ While this was not a one-shot sealed-bid auction, in that it permitted multiple bids in multiple rounds plus a separate set of sealed bids (as will be explained in more detail below), it used the VCG pricing rule, with a reservation price. As can be seen in the example, Swisscom won more blocks and paid more than 30% less than Sunrise. Orange paid the reservation price.

10.3.3 Budgets in VCG auctions

The VCG auction is particularly poorly suited to cases in which bidders face budget constraints. The following example has three bidders and three lots, and assumes budgets are binding for all bidders:

,,			
Bidder/bid for	1 lot	2 lots	3 lots
Bidder B1	100	100	100
Bidder B2	100	100	100
Bidder B3	100	100	100

Three-bidder, three-lot VCG auction

In this example, each bidder wins one lot and pays zero. Now, if one bidder

 $^{^6 \}rm One$ of the 1800 MHz blocks was 20 MHz, and the blocks were not all identical. However, Swisscom won better blocks than Sunrise.

Figure 10.1: Swiss 2012 Auction

Frequency			
band	Orange	Sunrise	Swisscom
800 MHz	20 MHz	20 MHz	20 MHz
900 MHz	10 MHz	30 MHz	30 MHz
1800 MHz	50 MHz	40 MHz	60 MHz
2.1 GHz	40 MHz	20 MHz	60 MHz
2.1 GHzTDD	0 MHz	0 MHz	0 MHz
2.6 GHz	40 MHz	50 MHz	40 MHz
2.6 GHz TDD	0 MHz	0 MHz	45 MHz
Amount paid			
(CHF)	I 54'702'000	481'720'000	359'846'000

has entered a bid of 101 for two or three lots and zero for one lot, then that bidder could win two lots for 100.

This auction has one type of unstable pure strategy equilibrium. This would have one bidder offering 200 (or more) for all three blocks, and the other bidders offering zero each. Of course, the other bidders could offer more than zero, and the bidder making the large offer would then want to revise its offer. But there will be no set of offers which constitute a stable (perfect) Nash equilibrium.⁷

To see this, suppose there were bids for each bidder constituting a Nash equilibrium. In this example, each bidder will want to allocate its budget among the three blocks. If each of two bidders allocates its budget of 100 to one block, then, as above, a rival can win two lots by offering 101 for two lots, and zero for other combinations. If only one bidder allocates its budget to one lot, then the other two bidders will be competing for the remaining two lots. And the only possibility is that they each allocate their budgets to one lot. A similar analysis would establish that there is no way for the three bidders to each allocate their budgets to the three lots so that no one bidder would not want to revise its allocation.⁸

The experience in many spectrum auctions indicates that even in very large auctions with large publicly traded firms, budgets can be constraining.⁹ So this example can be quite relevant.

10.4 Core-Selecting Auctions

This section addresses some of the issues raised above. First, it describes conditions under which the final prices are not so low that a losing bidder would be willing to offer more than the winners have to pay. When an auction has this property, then the auction outcome is said to be in the

⁷For a discussion of stable and perfect equilibria see [69] or [45].

 $^{^{8}}$ See ([67]) for an analysis of this type of auction with two bidders in a first-price sealed-bid auction. They find that there is an equilibrium in mixed strategies.

 $^{{}^{9}}See [16] and [65].$

core.¹⁰ More generally, the auction outcome is in the core when there is no group of bidders that can bid together and exclude other bidders, and that together would then each have a better outcome *which the auctioneer would also prefer*, i.e., would result in higher auction revenues.

Notice that, in the above examples, when an auction results in a non-core allocation, at least one bidder viewed the lots as complements. More specifically, the value of a second lot was higher than that of the first. In this case, bidders B1 and B2 each individually *blocked* B3. So, when B1's and B2's bids were each high enough to block B3, then B1 and B2 will each have a zero Vickrey price. So if the lots were also complements for B3, so that the value of the second lot was no higher than the first, the outcome would be in the core. B1 and B2 would each to pay the value of the first lot to B3. Given that the value of the first lot for B3 was more than the value of the second, what B1 and B2 pay in total (twice B3's value of the first lot) is necessarily more than B3's value of both lots. So, B3 no longer blocks the outcome, and the allocation is in the core.

More generally, it has been shown¹¹ that when all goods are mutual substitutes, the auction outcome will be in the core. Thus, the Vickrey auction outcome is in the core.¹² As the above discussion and examples demonstrate, this need not be the case when the goods are complements. The reason is that when one bidder's offers for a set of lots is larger than the sum of its offers for the individual lots, or for non-intersecting subsets, then two or more rivals can individually prevent that bidder from winning. So, if each of those rivals offers enough to block the bidder with complements, then their Vickrey prices will not reflect the value of the losing bid. This type of situation cannot arise when the lots are always substitutes for all bidders.

One additional property of core-selecting auctions is that the outcome does not change if bidders merge, or if a bidding consortium breaks up. The logic is fairly straightforward. Suppose some set of bidders could earn more bidding as one bidder than if they split up into separate bidders. More specifically, let S_1 and S_2 be disjoint subsets of the set of all bidders, N, and suppose S_2 consists of shills, so that $w(S_2) = 0$, where w(S) denotes the

¹⁰The notion of the core is discussed in [28]. Also, [21] discussed a solution which is essentially the core in the context of a simple exchange economy.

 $^{^{11}}See [8]$

 $^{^{12}\}mathrm{Also},$ the ascending proxy auction is in the core.

amount that the coalition $S \subset N$ can guaranty for its members when bidding together in the auction when bidding together. Suppose $w(S) > w(S_1)$. Note that efficiency of core allocations implies w(T) + w(S) = w(N) and $w(S_1) + w(S_2) + w(T) = w(N)$, where $T = N \setminus S$. This implies that $w(S_2) = 0$, i.e., the set S_2 of shills get nothing on their own. On the other hand, consider the merging of S_1 and S_2 into S. Then w(S) + w(T) = w(N). But, also, $w(S_1) + w(S_2) + w(T) = w(N)$. So $w(S_1) + w(S_2) = w(S)$. Thus, mergers never increase profits. Summarizing:

Proposition 3 ([54]). An efficient direct auction mechanism has the property that no bidder can ever earn more than its Vickrey payoff by disaggregating and bidding with shills if and only if it is a core-selecting auction mechanism.

10.5 Ascending Package Auctions and the Combinatorial Clock Auction

The combinatorial clock auction (CCA) is a form of ascending package auction that specifically addresses the low revenue of VCG auctions, and can lead to a core outcome. An ascending package auction is a simultaneous auction in which prices rise from round to round, as in an SMR auction, but all bids are package bids. When bidding is straightforward, as when using surplus-maximizing proxy bidders, [8] have shown that the final allocation of an ascending package auction is in the core. This section describes the CCA.

The Swiss auction described in the previous section was a CCA. The CCA has also been used for selling spectrum licenses in a number of other countries, including Austria, Denmark, Ireland, and the UK, and is planned for Australia and Canada.

The CCA differs from an SMR auction in a number of important respects:

• Before a CCA (as before an SMR auction), the available lots are assigned points, and bidders must specify how many points of initial eligibility they want. Applying for more points usually requires larger financial guarantees, which can provide some disincentive for applying the maximum number of points.

- In each round a bidder has to submit new bids that meet its activity target. In other words, the activity must meet the activity target. In most CCAs, a bidder's activity target is 100% of its eligibility. If a bidder's bids fall short of this target, its eligibility is reduced in all future rounds to the new, lower activity.¹³
- All bids are package bids. This is unlike an SMR auction, in which some lots bid in a round can become provisional winning bids.
- Any bid from any round can become a winning bid.
- The clock phase ends, as in SMR auctions, when there is no longer any excess demand for any product.
- The clock phase does not end the auction. Bidders can submit bids in a "supplementary round" of bidding. The bids in the supplementary round must improve, in some sense, on bids made during the clock rounds, and be consistent.¹⁴
- All bids are mutually exclusive. While a bidder can submit different package bids which do not overlap, a bidder can only win one bid. So, if a bidder wants to submit an offer for two different packages A and B, it might want to submit separate offers for A, B and A combined with B.
- At the end of the supplementary round, winners are determined by maximizing the sum of bidder offers across all feasible combinations of bids made during the clock rounds and during the supplementary round. The prices paid are the VCG prices except when these prices are not in the core in which case some rule is applied to adjust prices to ensure that they are in the core.¹⁵

 $^{^{13}}$ In some CCAs bidders can at times bid on earlier-round packages whose points exceed eligibility when the relative price of the larger package has decreased. See [7].

¹⁴A revealed preference type activity rule is used to ensure consistency.

¹⁵See [22] for a discussion of alternative core adjustments.

We have summarized the main elements of the CCA rules. CCAs have now been used in perhaps a dozen or more auctions. The rules differ across auctions, especially in regard to the core adjustment and to how improvement and consistency in the supplementary round are implemented. What is generally required is that the offers in the supplementary round satisfy a version of a *revealed preference* rule.

What follows is an example illustrating a revealed preference activity rule. Suppose, for example, a bidder was bidding for x units of product A and yunits of product B. Further suppose that during the clock phase this bidder always bid on two units of A and two units of B when the price of A was no more than Δ_1 larger than price of B, and never bid on more than one unit of A when the price differential reached $\Delta_2 > \Delta_1$, i.e.,

$$D(p_A, p_B) = \begin{cases} (2, 2) & \text{if } p_A - p_B \le \Delta_1 \\ (1, 3) \text{ or } (2, 2) & \text{if } \Delta_1 < p_A - p_B < \Delta_2 \\ (1, 3) & \text{if } p_A - p_B \ge \Delta_2 \end{cases}$$

In this case, the bidder has indicated that it would prefer the second lot of product A over a third lot of B when the price differential is less than Δ_1 , and would prefer the reverse when the price differential exceeds Δ_2 . Precisely where this bidder would switch is not necessarily revealed during the auction.

A revealed preference rule would require that supplementary round offers for the package (2,2) never exceed the offer for (1,3) by more than Δ_2 or less than Δ_1 . The implementation of the revealed preference rule, often called a *relative cap rule* or RCR, requires a number of decisions about how to resolve potential inconsistencies during the clock rounds, as is discussed below.

What follows is a brief review of some properties of the CCA, and a few examples from actual auctions.

10.5.1 Complexity

The CCA allows bidders to submit bids on any or all packages of lots available in the auction. The number of combinations can be astronomically large – if there are N objects for sale, then there are $2^N - 1$ possible combinations. Finding winners and determining prices requires choosing, from among all mutually consistent sets of offers from all the bidders (that is, offers for which total demand does not exceed supply), the set that maximizes the sum of the reported values. Further, solving for prices requires recalculating the optimal solution when each bidder is excluded. If the are n^j offers from bidder j, $j = 1, 2, \ldots, J$, then there are $n^1 \times n^j \times \cdots \times n^J$ possible combinations of bids to compare. Unless these numbers are very small, bidders cannot possibly submit bids for all combinations, nor can the auctioneer calculate winners and prices.

This complexity problem is of some practical concern. For example, in the Swiss auction described above, there were ten categories of licenses (across seven bands). There were over 60 million possible packages for each of the three bidders, and over 2×10^{23} possible combinations in all.¹⁶ Auctioneers are forced to limit the number of combinations, and bidders cannot fully express preferences. So, as a practical matter the CCA is not a useful mechanism for eliciting information from bidders about preferences.

Nor is the CCA the most efficient way of eliciting information about values. For example, suppose there are four products, and a bidder wants at least 1 unit of each product and a total of 6 units. The bidder may also want to purchase up to 4 units of each product. The CCA would require the bidder to submit values for 256 combinations. In practice the bidder may have a value for the base package of one unit of each product, and then an incremental value for each additional unit. If that is the case, the bidder would really need to submit only 13 values.

10.5.2 Bidder incentives

While the CCA is complex, it does provide bidders with some incentives to report true values – in that it uses Vickrey pricing. So, what follows examines where bidder incentives may differ from those in a standard Vickrey auction.¹⁷ The main differences between a VCG auction and a CCA arises

 $^{^{16}\}mathrm{There}$ were some limitations (e.g., spectrum caps) that reduced this number somewhat.

¹⁷The CCA differs from the ascending proxy auction of [8] in that there are no proxies in the CCA, and in the use of a second-price rather than a pay-as-bid rule.
when bidders payoffs depend not just on the price they pay but also on what rivals pay, and when there is a core adjustment.

Fairness

In the Swiss auction, Sunrise won fewer blocks and paid much more than Swisscom. Reportedly, Sunrise complained to the regulator. Unfortunately for Sunrise, the Swiss regulator, Bakom, had been made aware of this possibility and had indicated it was not relevant for its decision to proceed with the CCA.¹⁸ But a firm's decision-makers and their advisers are evaluated based on relative performance. Thus, inequitable outcomes may be economically inefficient when bidders' preferences include relative performance.

The first two large multi-product CCA auctions both resulted in some bidders paying more than rivals that won more spectrum. The discrepancy in the Swiss auction was quite large. There was also a smaller but quite direct discrepancy in the Dutch 4G auction (Figure 10.2). In that auction KPN won the same number of abstract blocks in the 800 MHz, 900 MHz and 1800 MHz bands as did Vodafone. However, KPN also won 25 MHz of additional spectrum in the 2600 MHz band. Yet, KPN's base price was €30 million *l*ess than Vodafone's. To avoid relatively poor performance in a CCA, bidders may need to deviate from bidding based purely on economic surplus as measured by their business models.

Bidding Long

In a CCA, as in Vickrey auctions, what rivals offer affects the amount any given bidder will pay. The CCA differs from a Vickrey auction in that there are multiple bidding rounds, and bidders are provided information about rivals between the rounds. The relative cap rule means that at the end of the clock phase a bidder may not have to offer full value to ensure what it wants in the auction. By not bidding full value, a bidder can improve its relative price. However, a bidder can be forced by a rival to bid more

 $^{^{18}}$ Bakom had been provided examples of this possibility in meetings I had with them prior to the auction.

Figure 10.2: Dutch 4G Auction Results

Dutch 4G Auction Results											
Band/Bidder	KPN	Vodafone	T-Mobile	Tele 2							
800 MHz	2 x 10 MHz	2 x 10 MHz		2 x 10 MHz							
900 MHz	2 x 10 MHz	2 x 10 MHz	2 x 15 MHz								
1800 MHz	2 x 20 MH	2 x 20 MH	2 x 30 MHz								
2100 MHz	2 x 5 MHz	2 x 5 MHz									
1900 MHz			I 4.6 MHz								
2600 MHz TDD	25 MHz		25 MHz								
Base Price	€ 1,349,851,000	€ 1,380,793,000	€ 910,582,000	€ 160,813,000							
Assignment round delta	€ 2,001,000	€ 7,000	€ 99,000								
Total price paid	€ 1,351,852,000	€ 1,380,800,000	€ 910,681,000	€ 160,813,000							

Round	Price/Bidder	Bidder 1 demand	Bidder 2 demand	Bidder 3 demand	B1 bid amount	B2 bid amount	B3 bid amoun
1	45	4	2	2	180	90	90
2	50	4	2	2	200	100	100
3	55	4	1	1	220	55	55
4	60	2	1	1	120	60	60
8	80	2	1	1	160	80	80

Table 10.5: Bid Long

for what it wants, when the rival keeps the clock phase from closing. The following example (Table 10.5.2) assumes three bidders seeking four lots.¹⁹

If the auction were to end after round 4, then, assuming no supplementaryround bids, bidder B1 would win two lots, and B2 and B3 would each win one. B1 would pay 80, because B2 and B2 had, in round 3, offered a combined 200 for the four lots, but only won one lot each, worth a combined 120. However, if the clock phase were to end four rounds later at a price of 80, then B1 would only pay 40. Note that B2 and B3 can offer up to 110 each in the supplementary round for two lots (as the maximum allowed offer is based on the price at which the bidder reduced eligibility). So, B1 may not necessarily benefit so much from extending the clock rounds, and will risk winning all the blocks. But there are clearly incentives to depart from straightforward bidding.

Budget Constraints

When bidders have fixed budgets, a CCA and a VCG present very challenging decision problems for the bidders. The first example is a variation of the chopsticks problem analyzed by [67].²⁰ Suppose there are two bidders and three lots. Further suppose each bidder has a budget of 1, and has to enter a separate bid for each of lots 1, 2, and 3.²¹ Suppose, too, that the amount a bidder pays is the amount the rival bidder would have paid for its lots. Further, suppose that the value of each lot exceeds a bidder's budget. If a

¹⁹This example was suggested by Jon Levin and Andy Skrzypacz. See also [1] and [37] for discussion of bidding strategy in CCAs.

²⁰Rosenthal and Szentes's [67] paper was motivated by Rosenthal's work on a sealed-bid spectrum auction planned in the Netherlands in 1998.

²¹If the bidders indicate an offer for one lot, two lots, or three lots, and the lots are identical, the analysis will differ, but there still won't be an equilibrium in pure strategies.

bidder's total offers ever exceeds its budget, the Vickrey price can also exceed its budget. So, if a bidder is not allowed to take any such risks, the bids for one or any combination of the three lots can never exceed 1 in total. But this means that each bidder is virtually certain to pay less than 1, and will have money left over. There will never be a pure strategy equilibrium in this game – the only possible equilibria involve randomization.

Now, consider the same bidders in a CCA. The bidders will tend to bid up to their budgets, and then drop. If bidders always behave in this manner, their winning offers and their maximum offers will equal their budgets. Then the Vickrey prices are always zero.²² In experiments and in practice, it does appear that bidders may have flexible budgets or take some risks. The data in actual auctions is confidential, and so it cannot be verified to what extent bidders risk exceeding budgets. Experience suggests, however, that such behavior does indeed occur

Experimental data suggest bidders will take very limited risks.²³ In a CCA, bidding above budget can be nearly riskless, as prior-round bids make the probability of those risky bids' winning essentially zero. Nevertheless, it does appear that bidders cannot easily assess risks, and they bid cautiously. The net effect is that payments will tend to be below budgets.

10.6 Comparing the Performance of the CCA and Other Auctions

There are now a large number of papers comparing the performance of the different auction formats.²⁴ What follows is a basic analysis of the differences. To compare a CCA and an SMR auction, suppose there are ten lots available, one large bidder that will purchase all ten lots, and two small bidders each seeking one lot.(Table 10.6)

If this auction were to be conducted as an SMR or clock auction, in which bidders paid their bid amounts, bidder B1 would win, but must pay a round

 $^{^{22}}$ [8] have also noted the decision problem posed by budget constraints.

 $^{^{23}}$ See [44].

 $^{^{24}}$ See [8] for some discussion. See also [12].

10.6. COMPARING THE PERFORMANCE OF THE CCA AND OTHER AUCTIONS185

Round	Price/Bidder	Bidder 1 demand	Bidder 2 demand	Bidder 3 demand	B1 bid amount	B2 bid amount	B3 bid amount
1	45	10	2	2	450	90	90
2	50	10	2	2	500	100	100
3	55	10	1	1	550	55	55
4	60	10	0	0	600	60	60

Table 10.6: Comparing CCA and SMR auctions

																				Opportunity	Base
Bidder	LA	LB	LC	LD	LE	LF	LG	LH	LI	LJ	LK	LL	LM	LN	LO	LP	LQ	Usage	Bid	cost	price
Qualcomm	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	Η	20,000	8,334	8,334
Bidder							Wit	nout	Qu	alco	omn	n						Usage	Bid		
ePortal								1	1	1	1	1	1	1				L	3,382		
Vectone			1	1	1													Н	2,338		
WorldSpace																	1	Н	2,614		
ePortal								1	1	1	1	1	1	1				L	3,382		
Vectone			1	1	1													н	2,338		
WorldSpace																	1	L	2,614		

Figure 10.3: UK L Band Auction

4 price of 60 for each lot, for a total of 600. On the other hand, in a CCA this bidder would have to pay only 200, or at most 220, for all ten lots. Thus, the CCA appears to provide an advantage to the large bidder. And if a large bidder wants to foreclose the small one, the costs of doing so may be much less in the CCA than in the SMR auction.

This is not just a theoretical example. One of the first CCA auctions was one for the UK L band (Figure 10.3) in May $2008.^{25}$ That auction included 17 lots: 16 small ones, and one large one, which was roughly equivalent to three small lots.

The last clock-round prices were £871,000 per lot, or more than £16,500,000 in aggregate. However, QUALCOMM, which won all the lots, paid £8,334,000 for all 17 lots. QUALCOMM's final-round bid was even higher, £20,000,000. Revenue equivalence does not hold in this case. Revenue equivalence does not tend to hold when bidders are bidding for multiple lots. To date, the theoretical, empirical, and experimental evidence on CCA vs SMR auction

 $^{^{25}\}mathrm{See}$ of com.org.uk for the details.

revenues is inconclusive.

10.7 Conclusion

This chapter has examined some forms of package bidding. Package bidding is potentially a significant practical concern for bidders. VCG auctions, a form of package bidding, have some unfortunate properties, though they do provide bidders incentives to bid truthfully, and they achieve efficient outcomes. These properties have appeared in actual auctions.

The problem of low revenues can be addressed through core selection auctions. However, predefined packages may be more appropriate in certain cases, as has been seen in some of the experimental literature.

One other type of package bidding that has been proposed is menu auctions.²⁶ In a menu auction bidders are asked to report values for different packages. For example, if bidders can win any number k of blocks with $\underline{k} \leq k \leq \overline{k}$, then a bid is an amount offered for each such possible k. This type of auction provides bidders incentives to bid true incremental values – even in a payas-bid, i.e., first-price, auction.

The reason that menu auctions have this property is that the bidding level determines the overall discount that a bidder applies, and then each bidder will have an incentive to only bid incremental values for incremental blocks. However, menu auctions may have multiple equilibria.²⁷ However, they present a faster and simpler sealed-bid approach to package bidding.²⁸

A version of the menu auction was used in the French 4G auctions (Figure 10.4). There were two auctions, one for the 800 MHz bands and one for

 $^{^{28}}$ The following example, from [11], illustrates this point. Suppose there are two bidders, A and B, and two blocks, X and Y.

Block	Value for Bidder A	E1 bid	E2 bid	E3 bid	E4 bid
X	6	5	0	7	0
Y	5	0	3	6	0
$Z = X \cup Y$	8	7	6	5	3
Ø	0	0	0	0	0

186

 $^{^{26}}$ See [11] Also, see [42] which discusses how similar ideas can be applied to other types of assets, such as financial ones with different ratings.

 $^{^{27}}$ See [11] .

Figure 10.4:

Bidders in French 4G menu auction													
					Total amount								
Band	Bouyges	Free	Orange	SFR	paid (€M)								
2600 MHz	€ 228	€ 272	€ 287	€ 150	€ 937								
Blocks won	2×15MH	2x20MH	2×20MH	2x15MHz									
800 MHz	€ 683		€ 891	€ 1,065	€ 2,639								
Blocks won	2×10MHz	nono	2x10MHz	2x10MHz									

the 2600 MHz bands. In each case bidders were asked to indicate demand curves for a range of 2 to 6 blocks, and for particular blocks in case they were not equivalent. The winning bids were based on the set of offers that maximized the sum of the total values (with some weights used to reflect other objectives of the auctioneer, such as coverage). Bidders were required to pay bid amounts.

Block	Value for Bidder A	E1 bid	E2 bid	E3 bid	E4 bid
X	5	0	2	7	0
Y	6	3	0	6	0
$Z = X \cup Y$	7	3	0	5	0
Ø	0	0	0	0	0

In this example, the strategies E1, E2, E3, and E4 are each equilibria. The revenues are respectively 7, 6, 5, and 3. The multiplicity arises from de facto coordination of the bids.

Chapter 11

Final Remarks

Summary

This chapter provides an overview of the main points explained in the preceding chapters. The intent is to provide some guidelines to what aspects of auction theory and experience may be relevant in any particular situation.

11.1 Introduction

This book is intended as a primer and not an encyclopedic guide or a survey of auction theory and practice. The intent is to provide some guidelines to what aspects of auction theory and experience may be relevant in any particular situation. This chapter highlights some of the key issues for the various types of auctions encountered in practice.

11.2 Game theory and bidder incentives

Any analysis, whether from the bidder or from the auctioneer perspective, needs to start with the incentives and constraints of the bidders. Forgetting to take account of incentives will result in surprises for both the bidder and the auctioneer. It is all too common for an auction design to fail in this way. The standard on-line English auction with a fixed closing rule creates large incentives for bidders to wait unit the last second. This makes the final outcome quite random.

In multi-attribute auctions, scoring rules can create incentives to distort values, leading to very perverse outcomes. This is what happened in the California QF auction, in which the suppliers were paying to supply energy, even during high-demand hours. In multi-object auctions, the structure of the auction can limit arbitrage opportunities or even present bidders with incentives to distort demand. There are many examples in which this occurred – the New Zealand spectrum auctions in the 1980s were one. But even the Google ad auctions for position limit bidders' ability to arbitrage price differentials for substitutes. This was also the case in the Mexican 2G auctions, in which activity rule parameters forced bidders to make sequential decisions; the result was that in some regions, larger and more valuable licenses sold for less than smaller licenses.

Budget constraints in multi-object auctions can confront bidders with impossible decision problems. For example, in a multi-object, sealed-bid auction, whether pay-as-bid or Vickrey, a bidder will be required to predict final prices in determining how to bid. This is part of what happened in the New Zealand auction for radio stations, cited above. But the auction design can be modified to make bidding decisions much simpler. For example, bidders can be allowed to submit demand schedules rather than separate bids for each object.

11.3 Bidder incentives and revenue equivalence

Perhaps the most fundamental insight of auction theory is the revenue equivalence theorem (RET). The form of the pricing rule affects bidder incentives – at times, in such a way that two seemingly quite different auctions can result in the same outcome. The debate about whether to use a pay-as-bid or a second-price auction often fails to take account of the RET. This is not to say that one can always rely on the RET to say that auction design and bidding strategy do not matter. The RET applies only in some situations, and it is more limited in its applicability to multi-object auctions than to single-object auctions. Moreover, even when the RET applies, the decision problem facing a bidder need not be the same in two different auctions that should result in the same outcome. For example, in a second-price auction for a single object, bidding value is a dominant strategy. The decision problem facing a bidder is relatively straightforward. However, even in the case of independent private values, where the RET does apply, a bidder in a first-price sealed-bid auction has to solve a complex optimization problem to determine its optimal bidding strategy. So, while theory would predict outcome equivalence, bidder decisions are more straightforward, and, as a result, it would appear that the outcome should be more predictable in a second-price auction than in a first-price auction.

11.4 Single-object vs multi-object auctions

Bidders face inherently more complex decisions in multi-object auctions than in single-object ones. When there are substitutes, bidders will need to decide what objects to bid more aggressively for. When there are complements, bidders face exposure problems. The form of the auction, including the nature of the bid, the pricing rule, and other provisions, has an enormous effect on the decision problem facing bidders, and therefore the auction outcome. Auctioneers want to carefully consider options for the auction format and its influence on bidding behavior as well.

One of the most significant successes in putting economic theory into practice has been the development of the SMR auction and other simultaneous auctions that simplify decision problems facing bidders. Simplifying decision problems can result in more efficient outcomes, higher revenues for the auctioneer, *and* more surplus for the bidders.

Package bidding, or combinatorial auctions, can be utilized for situations in which bidders face complementarities. In those cases in which there are known complementarities that are the same for many bidders, and little conflict over which packages are best, using predefined packages can be desirable in an SMR auction. However, this requires that the auctioneer have some knowledge about bidder preferences, and that is not always the case. The Vickrey pricing rule appears particularly unattractive for combinatorial auctions, especially when bidders have budget constraints. Package bidding can also introduce significant complexity into the bidding process. The complexity itself can result in more bidding errors and adversely affect the outcome.

11.5 Sequential auctions

Many times, auctions are not one-time events: similar or identical objects are auctioned at different times. This means that bidders in an early auction will want to forecast prices in deciding how to bid. One very central result of auction theory, the martingale theorem, states essentially that the expected price in one auction is the realized price in the previous auction. However, when bidders' opportunities and values can change over time, this result needs modification. In that case, the decision how to bid can matter. And the decision about how much to auction at each date can affect the outcome for the auctioneer.

Sequential auctions can occur for complements too. In this case, the early winner will tend to have an advantage in subsequent auctions. This is especially true in auctions for renewing franchises where there are complementary sunk and durable investments. In that case, prices in the first auctions will tend to be above values, in subsequent auctions below values; they will not provide much direct information about true underlying values.

11.6 Bidding behavior

Most of the analysis in this book has assumed, explicitly or implicitly, that bidders act rationally to maximize their expected payoffs. This assumption is clearly unrealistic in many situations, even in very large auctions with very sophisticated bidders. However, starting with this benchmark is also useful.

Perhaps most useful is working through the decision problems facing bidders. Referring again to the comparison between first-price and second-price auctions, it is clear that bidding strategy is generally simpler in a secondprice auction – at least for single-object auctions. This is not to say that behavior does not matter. In experimental settings bidders at times can systematically depart from rational play.¹ At other times, the outcome may not be so sensitive to deviations.² Where there is a lot at stake, bidders have time to analyze the decisions, and the decision problems are straightforward, then the assumption of rational bidding is likely to be most valid. However, it should also be clear from this primer that an analysis of the decision problem facing a rational bidder can also be useful even where bidders lack the ability, resources, incentives, or time to come up with optimal bidding strategies. Indeed, at times the decision costs of coming up with an optimal decision can far outweigh the potential benefits.

When a bidder faces very random outcomes, then the decisions are especially difficult. Some auctions, which rely on outguessing rivals, have intrinsic randomness, in that they only have mixed equilibria. In some cases it is possible to design an auction to avoid this randomness. This type of uncertainty can, in most cases, only result in worse expected outcomes for both the bidders and the auctioneer.

11.7 Optimal auctions

A fundamental contribution to auction theory is Myerson's [57] (1981) paper "Optimal Auction Design." Auctions inherently involve some informational asymmetries: bidders know their own values, which are unknown to the auctioneer and possibly to each other. While bidders may end up revealing values in any optimal auction, this does not mean that when asked to report values, bidders will necessarily do so. Moreover, even in Vickrey auctions, in which bidders have strong incentives to report true values, the auctioneer will not necessarily maximize its expected revenues. Indeed, the optimal auction from a seller's perspective will necessarily include a reservation price.

Thus, the optimal auction – for example, the one that maximizes seller revenues – will necessarily leave some surplus for the bidder. This is due to the

 $^{^1\}mathrm{See}\ [39]$.

 $^{^{2}}$ See [14] .

underlying informational asymmetries – asymmetries that leave information rents to the bidders. Thus, the goals of eliciting truthful valuations from bidders and maximizing auction proceeds can conflict, even if bidders always end up bidding in a way that can be used to infer true values.

11.8 Auction design, management, and strategy

Putting the principles described in this primer to use involves many decisions. These decisions can be very complex. The goal in this primer is not to provide a checklist to allow a bidder or an auctioneer to decide on the best bidding strategy or best auction format. Rather, this primer is intended to provide the bidder or the auctioneer with a way of starting an analysis of its decision problem. The basic tools of game theory, in Chapter 2, should be applied. And the analysis of auction design or strategy will depend on whether there is a single object or multiple objects (which can be substitutes or complements), and on whether each bid will include a single lot or multiple lots. However, good design and management decisions and good bidding strategies will almost always require some analysis based on the tools given, though no formula can be applied in all situations. And even where formulas can be applied, care is needed to avoid misuse.

11.9 Participation

As a final remark, auctions don't work without bidders. Not only must there be bidders, the bidders need to submit competitive and meaningful bids. Auction rules and management can make bidder decisions difficult, which discourages bidders from making meaningful bids. This is true in even simple auctions, such as simultaneous sealed-bids. But, it is also true in more complex auctions, such as Combinatorial Clock Auctions, bidders can be confronted with impossibly difficult decisions. Auctions can fail if bidders fail to be informed of the opportunity, or if bidder qualification criteria or payment rules are too lax.³

11.10 Auctions and markets

This Primer is intended to illustrate some pitfalls, but more important explain how bidder incentives need to be factored into any analysis of auction design and bid strategy. This Primer explains how the auction design can affect bidder incentives, which, in turn affect the auction outcome.

Auctions are a form of market mechanism. Auctions differ from other types of markets in that they have formal rules governing offers, allocations and pricing. An auction is particularly well suited for one-off or intermittent transactions with relatively few bidders. Higher frequency transactions typically occur in other types of markets and exchanges.

That auction design, management and strategy can have a large impact on the final allocation and prices suggests that auctions are not reliable nor predictable. This Primer should make it clear, however, that auctions, if managed well, can be a useful and efficient type of market mechanism for pricing and making allocation decisions.

 $^{^{3}{\}rm The}$ Nextwave bankruptcy tied up a significant fraction of the available cellular frequency for over five years. See FCC V. NEXTWAVE PERSONAL COMMUNICATIONS INC. (01-653) 537 U.S. 293 (2003) 254 F.3d 130.

196

Bibliography

[1]

- [2] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society*, pages 1041–1063, 1990.
- [3] Blaise Allaz and Jean-Luc Vila. Cournot competition, forward markets and efficiency. *Journal of Economic theory*, 59(1):1–16, 1993.
- [4] Orley Ashenfelter. How auctions work for wine and art. The Journal of Economic Perspectives, 3(3):23-36, 1989.
- [5] Orley Ashenfelter and Kathryn Graddy. Auctions and the price of art. Journal of Economic Literature, 41(3):763–787, 2003.
- [6] Lawrence M Ausubel and Peter Cramton. Virtual power plant auctions. Utilities Policy, 18(4):201–208, 2010.
- [7] Lawrence M Ausubel and Peter Cramton. Activity rules for the combinatorial clock auction. Department of Economics, University of Maryland, Discussion Paper, 2011.
- [8] Lawrence M Ausubel and Paul Milgrom. Ascending auctions with package bidding. Frontiers of theoretical economics, 1(1):1–42, 2002.
- [9] Ross Baldick, Ryan Grant, and Edward Kahn. Theory and application of linear supply function equilibrium in electricity markets. *Journal of Regulatory Economics*, 25(2):143–167, 2004.
- [10] Dan Bernhardt and David Scoones. A note on sequential auctions. *The American Economic Review*, 84(3):653–657, 1994.

- [11] B Douglas Bernheim and Michael D Whinston. Menu auctions, resource allocation, and economic influence. *The quarterly journal of economics*, 101(1):1–31, 1986.
- [12] Martin Bichler, Pasha Shabalin, and Jurgen Wolf. Efficiency, auctioneer revenue, and bidding behavior in the combinatorial clock auction. Technical report, Working Paper, TU Munchen, 2011.
- [13] Severin Borenstein, James B Bushnell, and Frank A Wolak. Measuring market inefficiencies in california's restructured wholesale electricity market. *American Economic Review*, pages 1376–1405, 2002.
- [14] James N Brown and Robert W Rosenthal. Testing the minimax hypothesis: A re-examination of o'neill's game experiment. *Econometrica: Journal of the Econometric Society*, pages 1065–1081, 1990.
- [15] Sandro Brusco and Giuseppe Lopomo. Collusion via signalling in simultaneous ascending bid auctions with heterogeneous objects, with and without complementarities. *The Review of Economic Studies*, 69(2):407– 436, 2002.
- [16] Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. Technical report, National Bureau of Economic Research, 2009.
- [17] Edward C Capen, Robert V Clapp, and William M Campbell. Competitive bidding in high-risk situations. *Journal of petroleum technology*, 23(6):641–653, 1971.
- [18] Peter Cramton and Jesse A Schwartz. Collusive bidding: Lessons from the fcc spectrum auctions. *Journal of regulatory Economics*, 17(3):229– 252, 2000.
- [19] R Day and P Milgrom. forthcoming. optimal incentives in core-selecting auctions. handbook of market design, 2010.
- [20] Robert Day and Paul Milgrom. Core-selecting package auctions. international Journal of game Theory, 36(3-4):393–407, 2008.
- [21] Francis Ysidro Edgeworth. Mathematical psychics: An essay on the application of mathematics to the moral sciences. Number 10. CK Paul, 1881.

- [22] Aytek Erdil and Paul Klemperer. A new payment rule for core-selecting package auctions. Journal of the European Economic Association, 8(2-3):537-547, 2010.
- [23] Jun Byoung Heon Fan, Cuihong and Elmar Wolfstetter. Optimal information disclosure in a license auction with downstream interaction. Technical report, Working Paper, Korea University, 2013.
- [24] Ronald A Fisher. On the mathematical foundations of theoretical statistics. Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, 222:309– 368, 1922.
- [25] James W Friedman. Cooperative equilibria in finite horizon noncooperative supergames. *Journal of Economic Theory*, 35(2):390–398, 1985.
- [26] Peter M Garber. Tulipmania. Journal of Political Economy, pages 535– 560, 1989.
- [27] Richard J Gilbert and David MG Newbery. Preemptive patenting and the persistence of monopoly. *The American Economic Review*, pages 514–526, 1982.
- [28] DB Gilles. Solutions to general non-zero-sum games. contributions to the theory of games vol. iv. Annals of Math Studies, 40:47–85, 1959.
- [29] Jacob K Goeree and Charles A Holt. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior*, 70(1):146–169, 2010.
- [30] Edward J Green and Robert H Porter. Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society*, pages 87–100, 1984.
- [31] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica: Journal of the Econometric Society*, pages 427–438, 1977.
- [32] Richard J Green and David M Newbery. Competition in the british electricity spot market. *Journal of political economy*, pages 929–953, 1992.

- [33] Paul R Gribik. Designing an auction for qf generation resources in california: What went wrong? The Electricity Journal, 8(3):14–23, 1995.
- [34] Veronika Grimm, Frank Riedel, and Elmar Wolfstetter. Low price equilibrium in multi-unit auctions: the gsm spectrum auction in germany. International journal of industrial organization, 21(10):1557–1569, 2003.
- [35] Kenneth Hendricks and Robert H Porter. An empirical study of an auction with asymmetric information. *The American Economic Review*, pages 865–883, 1988.
- [36] Ali Hortacsu and Steven L Puller. Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics*, 39(1):86–114, 2008.
- [37] Vladimir Janssen, Maarten Karamychev. Gaming in combinatorial clock auctions. 2013.
- [38] Paul Joskow and Edward Kahn. A quantitative analysis of pricing behavior in california's wholesale electricity market during summer 2000. In *Power Engineering Society Summer Meeting*, 2001, volume 1, pages 392–394. IEEE, 2001.
- [39] John H Kagel and Dan Levin. Auctions: A survey of experimental research, 1995-2008. *Handbook of experimental economics*, 2, 2008.
- [40] Alfred E Kahn, Peter C Cramton, Robert H Porter, and Richard D Tabors. Uniform pricing or pay-as-bid pricing: a dilemma for california and beyond. *The Electricity Journal*, 14(6):70–79, 2001.
- [41] Paul Klemperer. Auctions: Theory and practice. *Economics Group*, *Nuffield College, University of Oxford, Economics Papers*, 2004.
- [42] Paul Klemperer. The product-mix auction: A new auction design for differentiated goods. Journal of the European Economic Association, 8(2-3):526-536, 2010.
- [43] Paul D Klemperer and Margaret A Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica: Journal of the Econometric Society*, pages 1243–1277, 1989.

- [44] Stephan Knapek and Achim Wambach. Strategic complexities in the combinatorial clock auction. 2012.
- [45] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica: Journal of the Econometric Society*, pages 1003–1037, 1986.
- [46] Vijay Krishna. Auction theory. Academic press, 2009.
- [47] Evan R Kwerel and Gregory L Rosston. An insiders' view of fcc spectrum auctions. *Journal of Regulatory Economics*, 17(3):253–289, 2000.
- [48] Nathan Larson and David Salant. Equilibrium in wholesale electricity markets. *Energy Research Center of the Netherlands (ECN)*, 2003.
- [49] Colin Loxley and David Salant. Default service auctions. Journal of Regulatory Economics, 26(2):201–229, 2004.
- [50] Robert C Marshall and Leslie M Marx. Bidder collusion. Journal of Economic Theory, 133(1):374–402, 2007.
- [51] David McAdams. Adjustable supply in uniform price auctions: Noncommitment as a strategic tool. *Economics Letters*, 95(1):48–53, 2007.
- [52] R Preston McAfee and Daniel Vincent. The declining price anomaly. Journal of Economic Theory, 60(1):191–212, 1993.
- [53] Paul Milgrom. Procuring universal service: Putting auction theory to work. Lecture at the Royal Swedish Academy of Sciences, 9, 1996.
- [54] Paul Milgrom. *Putting auction theory to work*. Cambridge University Press, 2004.
- [55] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.
- [56] Milton Mueller. New zealand's revolution in spectrum management. Information Economics and Policy, 5(2):159–177, 1993.
- [57] Roger B Myerson. Optimal auction design. Mathematics of operations research, 6(1):58–73, 1981.

- [58] Harry J Paarsch and Han Hong. An introduction to the structural econometrics of auction data. *MIT Press Books*, 1, 2006.
- [59] David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. Combinatorial auction design. Proceedings of the National Academy of Sciences, 100(19):11153–11157, 2003.
- [60] Robert H Porter. A study of cartel stability: The joint executive committee, 1880-1886. The Bell Journal of Economics, pages 301–314, 1983.
- [61] John G Riley and William F Samuelson. Optimal auctions. The American Economic Review, 71(3):381–392, 1981.
- [62] Michael H Riordan and David J Salant. Preemptive adoptions of an emerging technology. *The journal of industrial economics*, pages 247– 261, 1994.
- [63] David Salant. Auctions and regulation: Reengineering of regulatory mechanisms^{**}. Journal of Regulatory Economics, 17(3):195–204, 2000.
- [64] David Salant. Multi-lot auctions. In Obtaining the Best from Regulation and Competition, pages 41–64. Springer, 2005.
- [65] David J Salant. Up in the air: Gte's experience in the mta auction for personal communication services licenses. Journal of Economics & Management Strategy, 6(3):549–572, 1997.
- [66] Paul A Samuelson. A note on the pure theory of consumer's behaviour. Economica, 5(17):61–71, 1938.
- [67] Balázs Szentes and Robert W Rosenthal. Three-object two-bidder simultaneous auctions: chopsticks and tetrahedra. *Games and Economic Behavior*, 44(1):114–133, 2003.
- [68] Jean Tirole. The Theory of Industrial Organization: Jean Tirole. MIT press, 1988.
- [69] Eric Van Damme. Stability and perfection of Nash equilibria. Springer Verlag, 1991.

- [70] Gerard J Van den Berg, Jan C Van Ours, and Menno P Pradhan. The declining price anomaly in dutch dutch rose auctions. *The American Economic Review*, 91(4):1055–1062, 2001.
- [71] Hal R Varian. Position auctions. International Journal of Industrial Organization, 25(6):1163–1178, 2007.
- [72] John Vickers. The evolution of market structure when there is a sequence of innovations. *The Journal of Industrial Economics*, pages 1–12, 1986.
- [73] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. The Journal of finance, 16(1):8–37, 1961.
- [74] Leon Walras and William Jaffé. Elements of pure economics; or, The theory of social wealth. American Economic Association and the Royal Economic Society, 1954.
- [75] Frank A Wolak. Measuring unilateral market power in wholesale electricity markets: the california market, 1998-2000. The American economic review, 93(2):425–430, 2003.
- [76] Elmar Wolfstetter. Auctions: an introduction. Journal of Economic Surveys, 10(4):367–420, 1996.