# 1 Introduction

Market design is a kind of economic engineering, utilizing laboratory research, game theory, algorithms, simulations, and more. Its challenges inspire us to rethink longstanding fundamentals of economic theory.

Paul Milgrom, 2008

The digital economy has led to many new services where supply is matched with demand for various types of goods and services. While only a few years ago academics were mainly concerned with models describing markets, more and more people and organizations are now in a position to design market rules that are being implemented in software. The design of markets is challenging as it is necessary to consider the strategic behavior of market participants, psychological and cognitive factors, and computational problems in order to implement the objectives of a designer. Market models in economics have not lost their importance, but recent years have led to many new insights and principles for the design of markets which are beyond traditional economic theory. In this book, we study *market design*, an engineering field at the intersection of computer science, economics, information systems, and operations research concerned with the design of real-world markets.

Consider a transportation market with multiple lanes on a transportation network, one shipper and multiple carriers. The shipper has a set of truckloads to be shipped from different origins to various destinations. The carriers are available to meet the transportation demand, and they are invited by the shipper to submit sealed bids. Carriers typically have preferences for bundles of lanes on a route such that there are no empty lanes without a shipment, and they also submit respective bids on packages of lanes. This simple logistics procurement example leads to challenges for the shipper.

First, the shipper needs to determine an optimal allocation of bundle bids such that his costs are minimized. Cost minimization is a computationally hard optimization problem in this setting. Second, carriers want to maximize their payoff and it is not necessarily in their interest to reveal their costs truthfully. They might charge a high markup on some lanes where they expect no competition or they might not bid on other lanes where they expect high competition. However, strategic manipulation of this sort can lead to suboptimal allocations of lanes to carriers, high procurement costs for the shipper, and high bid preparation costs for the carriers, who would benefit from information about their competitors. Ideally, a shipper would have an economic mechanism where carriers have incentives to reveal their costs truthfully, and he or she can then determine the cost-minimal allocation optimally. In sales auctions the objectives are typically the maximization of auctioneer revenue or overall welfare. The latter is also referred to as (allocative) efficiency.

Overall, there are fundamental strategic and computational challenges in the design of such multi-object auctions, which are at the core of market design. The 2012 Nobel Memorial Prize in Economic Sciences for Alvin Roth and Lloyd Shapley honored research in this field. Market design uses economic theory, mathematical optimization, systems design, experiments, and empirical analysis to design market rules and institutions. Fundamentally, it asks how the design of the rules and regulations of a market affects the functioning and outcomes of that market. The study includes auction markets but also markets without money such as matching markets, which are used in the assignment of students to courses, in school choice programs, and in kidney exchanges.

*This textbook focuses on the design and analysis of efficient multi-object market mechanisms.*

## 1.1 Market Design and Mechanism Design

Market design has theoretical foundations in *mechanism design theory*, an analytical framework for thinking about what a given economic institution can achieve when the information necessary to make decisions is dispersed and privately held. The 2007 Nobel Memorial Prize in Economic Sciences to Leonid Hurwicz, Eric S. Maskin, and Roger B. Myerson was awarded in this field of economic theory. Mechanism design uses an axiomatic method, deriving results from a number of basic assumptions about utility functions or overall design goals of a mechanism. For example, mechanism design has been used to characterize the utility functions of market participants and mechanisms for which the truthful revelation of preferences is an equilibrium, i.e., a situation from which no participant wants to deviate (aka incentive-compatible mechanisms). The mechanism design literature shows that environments which allow for welfare maximization such that participants cannot make a loss and have strong incentives to bid truthfully are limited and require strong assumptions. The celebrated Vickrey–Clarke–Groves mechanism provides dominant strategies for agents to reveal valuations truthfully if they have independent and private values and they can maximize their payoff. If these assumptions are relaxed, so that bidders can have private budget constraints or valuations that are not independent, such positive results no longer exist (unless preferences can be characterized by a single parameter only). Overall, environments where truthful bidding satisfies a strong equilibrium solution concept are quite restricted. Often neither the bidders' preferences in real markets satisfy the assumptions necessary for truthful mechanisms nor is it possible to implement the required mechanisms for practical reasons. Actually, the Vickrey–Clarke–Groves mechanism is rarely used in the field. Still, it is important to understand the assumptions and mechanisms which would make truthful bidding an equilibrium bidding strategy.

*Market design* starts from the requirements in the field, eliciting the preferences and financial constraints of participants (which might be different from pure payoff maximization), information about competitors available to participants, objectives, and constraints on the allocation problem and on the award process. In many cases these requirements differ from those assumed in mechanism design theory, demanding the development of new models that capture the specifics of a certain type of market. Examples are models of markets for procurement and logistics, for spectrum sales, or for display advertising. Often the specifics of these markets might not allow for mechanisms which are incentive-compatible according to the established equilibrium solution concepts. So, market design takes the market environment as given and derives designs that satisfy some design goals (such as stability or budget balance of the outcome), while relaxing others. In other words, it analyzes tradeoffs and aims for "satisficing" solutions (Simon, 1996) to real types of market design problems rather than finding sets of assumptions that allow for "optimal" designs according to some design desiderata.

Market design complements mechanism design in the level of detail that is specified and analyzed. For example, in mechanism design valuations are typically described as an abstract function of the objects to be allocated, while the parametric forms of the valuation or cost functions, of the utility functions, and of the corresponding bid languages play a significant role in market design. Moreover, market designs almost always need to consider allocation constraints such as restrictions on the number of winners or the quantity awarded to individual bidders or groups of bidders. Therefore market design typically starts out as a mathematical programming task concerned with the design of an appropriate objective function, of appropriate constraints for the overall allocation, and of activity rules and with the definition of a bid language that lets bidders express their preferences in a succinct way. Milgrom (2017) provides an excellent and up-to-date view on market design and how it relates to economic theory.

## 1.2 Market Design and Mathematical Optimization

One thesis of this book is that *mathematical optimization*, in particular linear and integer linear programming, plays a central role in the design and analysis of multi-object markets and also provides a central foundation for theoretical models of multi-object markets. Many techniques and models introduced in this book use the theory of linear and integer linear programming in one way or another. Compared with algorithms designed for specific problems, there is an amazing versatility in linear and integer linear programming, which makes these techniques very important for market design. Algorithms for the solution of integer linear programs have seen substantial progress in the past two decades, allowing for totally new types of markets in different domains. Before that, in 1975 Leonid Kantorovich and Tjalling Koopmans had already received the Nobel Memorial Prize in Economic Sciences for their work on linear programming and optimal resource allocation. However, the advances in recent years have been staggering and have made many new markets possible that were not considered tractable just a few years ago. The subtitle of this book is reminiscent of the seminal book *Mechanism Design: A Linear Programming Approach* by Vohra (2011), who emphasized the role of linear programming in mechanism design theory.

However, market design is more than mathematical programming, as it takes into consideration bidding strategies and human behavior. Ignoring bidding strategies and different types of manipulation in market-based resource allocation is like optimizing a problem with the wrong parameters. As indicated earlier, strong game-theoretical solution concepts such as dominant-strategy or ex post incentive compatibility are desirable but they might not always be attainable, as a number of impossibility results in mechanism design show. These properties rest on assumptions about bidders' utility functions, which are often not given in the field. Still, in many markets it is possible to devise market rules for which manipulation becomes hard, given the uncertainties about other bidders and their valuations.

Overall, market design aims at design principles and rules for market institutions which are *robust against strategic manipulation and allow bidders to express their preferences*, so that the designer can aim for good or even optimal allocations. This is similar to network security, where designers aim for secure network protocols that are hard to tamper with, knowing that there is no absolute security. Note that "optimality" often refers to social welfare maximization, but the market designer might have different objectives or policy goals. In this sense, market design extends mathematical programming to situations with multiple decision makers. It is a *design science* in the sense of Herb Simon who received both the Nobel Memorial Prize in Economic Sciences (1978) and the Turing Award in Computer Science (1975). Simon's book *The Sciences of the Artificial* (Simon, 1996) motivated the development of systematic and formalized design methodologies relevant to many design disciplines, for example architecture, engineering, urban planning, medicine, computer science, and management science. Market design is therefore a very suitable name for the study of principles and methods in designing markets in the spirit of design science.

## 1.3 Outline of the Book

This textbook is intended to provide a one-semester course on market design. I am primarily targeting students with a background in computer science, information systems, mathematics, and management science. Hence, I will first introduce in Part I necessary concepts from game theory and mechanism design, as these students typically have not had the respective introductory courses. Parts II and III cover material which is more recent and is often not covered in microeconomics.

One prerequisite for this book is a familiarity with *linear and integer linear programming* and an introductory course in calculus and probability theory. There are many introductory textbooks on these subjects which would provide an excellent start for the topics discussed throughout this book. The appendices summarize important results from mathematical optimization and should serve as a convenient reference for the reader and a brief introduction for those who have not studied the respective courses.

As outlined earlier, this book focuses on the allocation of multiple objects with distributed decision makers. While single-object auctions are fairly well understood, multi-object auctions provide many more design options and they are more challenging to design and analyze. The goal of this book is to introduce important models and principles from game theory, mechanism design, and single-object auction theory for which there are several excellent textbooks. After basic concepts have been introduced, the book focuses on multi-object markets and their properties.

This leads to the three parts of the book: *Part I (Microeconomic Fundamentals)* discusses basic concepts from game theory, mechanism design, and single-object auction theory. The goal of this part is not to discuss each of these topics in breadth but to introduce the terminology and theory required for Parts II and III. The above topics are typically not taught in computer science and management science, and this makes it hard for students from these fields to understand the literature in market design. *Part II (Multi-Object Auction Design)* introduces the more recent theory and specific designs for multi-object markets. *Part III (Approximation and Matching Markets)* analyzes approximation mechanisms which have been designed for markets where the designer cannot hope to solve the allocation problem optimally. Then the design of matching markets where monetary transfers are not possible is discussed. Approximation and randomization play an important role in the recent matching literature. *Part IV (Appendices: Mathematical Optimization)* summarizes the main results from linear and integer linear optimization, on which I draw in Parts II and III.

There is always a tradeoff between the breadth and the depth of a textbook, and it is a matter of choice how much space is devoted to each single topic. If all the various literature streams in computer science, economics, and operations research were covered then a book on market design could easily cover thousands of pages, and it would clearly be beyond a one-semester course. This book provides a fair amount of detail in Part II and also in Part III, but Parts I and IV are limited to a necessary minimum for readers who have not taken the respective introductory courses. Selected references to the original literature are included in the text for those interested in learning more about a specific topic and in reading the original literature.

### 1.3.1 Part I Microeconomic Fundamentals

Chapter 2 introduces basic game-theoretical notions and solution concepts relevant for market design problems. Solution concepts for normal-form, extensive-form, and Bayesian games are introduced. There are many excellent textbooks on game theory. For example, Shoham and Leyton-Brown (2011) provides a comprehensive introduction for computer scientists, while I cover only a subset of the topics relevant for later chapters.

In Chapter 3 I discuss mechanism design theory, also known as inverse game theory. While the rules of a game are given in game theory, mechanism design tries to design rules such that certain goals and solution concepts are achieved. Market design problems can be considered as games where participants should have incentives to reveal their preferences for objects truthfully. If participants were truthful then the market designer would only need to solve an optimal allocation problem. The first section shows that aggregating general and unrestricted preferences is hard and that simple truthful mechanisms for general preferences are impossible. Then preferences are restricted to independent and private valuations and quasi-linear utility functions, where agents maximize payoff and monetary transfers are allowed. This restriction of preferences allows for a truthful mechanism with a dominant-strategy equilibrium. This means that bidders do not need information about competitors to bid optimally. The Vickrey–Clarke–Groves (VCG) mechanism is a celebrated result which is pivotal for much of the later discussion on auction design.

Chapter 4 introduces the traditional literature on single-object auctions. This literature provides a consistent framework for thinking about phenomena in auction markets, and it formalizes various assumptions. Equilibrium bidding strategies are discussed in a number of standard auction formats, before the outcomes of these auctions in equilibrium are compared. Krishna (2009) provides an excellent textbook which covers auction theory with a focus on Bayesian equilibrium analysis. This chapter is limited to key concepts providing a framework for the analysis of multi-object auctions later on.

### 1.3.2 Part II Multi-Object Auction Design

Chapter 5 introduces a taxonomy of multi-object auctions. *Multi-unit* auctions (for homogeneous objects) and *multi-item* auctions (for heterogeneous objects) are distinguished. For both types of auction, *sealed-bid* and *open* auction formats have been designed. Open auctions include those that are continuous or organized iteratively in rounds. They can be ascending or descending but always reveal some information about competitors throughout the process. Multi-item auctions are ubiquitous in industrial procurement, in logistics, and also in the public sector (Sandholm, 2012). Much of the remaining discussion is focused on multi-item auctions.

Chapter 6 is devoted to a simple and widespread multi-item auction format, the simultaneous multi-round auction (SMRA). It is a practical auction design and provides an excellent way to introduce some problems that arise in multi-item auctions. The strategic problems in SMRAs led to the development of combinatorial auctions. Combinatorial auctions allow bidders to express all types of preferences including complements and substitutes and can therefore be considered as the most general class of multi-object auction designs.

Chapter 7 introduces sealed-bid combinatorial auctions. The winner determination problem and the respective bid languages are discussed, before the various different payment rules are introduced. The VCG mechanism also leads to dominant-strategy equilibria in combinatorial auctions with payoff-maximizing bidders. However, in contrast with single-object auctions, this mechanism does not necessarily yield a competitive equilibrium, i.e., a stable solution where nobody wants to deviate. Also, the number of possible packages grows exponentially, so that a VCG mechanism might not be practical. The chapter introduces alternative payment rules and illustrates domain-specific bid languages addressing these issues. Combinatorial double auctions are also discussed, but the literature in this field is scarce.

Chapter 8 focuses on open and iterative combinatorial auctions, where bidders are able to outbid each other much as in an English auction. Assignment markets are a restricted type of market, where each bidder wants to win at most one out of several items. This restricted environment allows for ascending auctions with dominant strategies, and, what is more, they provide an excellent environment to illustrate algorithmic models of auctions. These types of model have received much recent attention and complement the literature on sealed-bid auctions. Also, in open multi-item auctions it is even more important to understand the interactions of bidders than it is in the single-item counterparts of these auctions. Note that most high-stakes multi-item auctions are open in format. Unfortunately, open combinatorial auctions cannot always have dominant strategies if modeled as a Bayesian game, where bidders have only distributional information about their competitors.

In Chapter 9 I discuss the single-stage and two-stage combinatorial clock auction formats. These auction formats have been used for the sale of spectrum licenses by governments around the world and provide an excellent example of practical market designs addressing the complexities of multi-object auctions. I also consider some pitfalls and problems that can arise in these auction formats.

### 1.3.3 Part III Approximation and Matching Markets

Chapter 10 analyzes market design problems in situations where the allocation problem is not tractable. Actually, the allocation problem of many real-world market design problems can be described as a combinatorial optimization problem and so computationally hard to solve. Unfortunately, the VCG mechanism provides dominant strategies only if the allocation problem can be solved exactly. Approximation mechanisms solve the allocation problem in polynomial time to give the objective function value of the underlying allocation problem only within a certain approximation ratio, but the mechanism provides incentives for the truthful revelation of valuations.

Finally, Chapter 11 provides an introduction to matching markets, which do not allow for monetary transfers. The assignment of students to courses at universities serves as a widespread example. There are environments where both sides of the market have preferences (e.g., students and course organizers) and others where only one side has preferences (e.g., the students). Interestingly, there are mechanisms, for both types of matching markets, where the truthful revelation of preferences is a dominant strategy for one side of the market. For example, the well-known deferred acceptance algorithm by Gale and Shapley provides dominant strategies for one side (e.g., the students) in a two-sided matching market, and the outcomes are stable, i.e., in our course allocation example there is no pair of students and course organizers who would want to switch. There are many connections to the theory of multi-object auction markets discussed in Chapter 8 and to approximation mechanisms as they are described in Chapter 10.

Chapter 12 provides an outlook on current trends and frontiers in market design.

### 1.3.4 Part IV Appendices: Mathematical Optimization

The appendices are intended as a reference to key results in *linear and integer linear programming* that are relevant to multi-object market design in Parts II and III of the book. Linear and integer linear programming is a vast field. I have assumed that readers have had an introduction to linear programming, and the appendices should serve only as a reference or refresher. A textbook introduction to linear programming can be found in Chvatal (1983) and to integer linear programming in Nemhauser and Wolsey (1988).

In the book I will often take a perspective including the reader and therefore will use "we" instead of "I" throughout.

## 1.4 Acknowledgements

Some sections of the book overlap with parts of papers that have been co-authored with students and colleagues, but I have combined them using a uniform notation with examples geared towards a textbook-style introduction to the field. These papers include Bichler (2010), Bichler *et al.* (2011a, 2013a, 2017b), Diebold *et al.* (2014), Kroemer *et al.* (2014). In preparing this text, I have benefitted from support and advice from many sources. I would like to express my gratitude to my current and former students Franz Diebold, Salman Fadaei, Andor Goetzendorff, Zhen Hao, Dennis Kraft, Richard Littmann, Stefan Mayer, Sören Merting, Per Paulsen, Ioannis Petrakis, Alexander Pikovsky, Stefan Schneider, Thomas Setzer, Pasha Shabalin, Florian Stallmann, Stefan Waldherr, Bernhard Waltl, Jürgen Wolf, and Georg Ziegler. I am thankful to co-authors and colleagues such as Gedas Adomavicius, Larry Ausubel, Haris Aziz, Oleg Baranov, Felix Brandt, Peter Cramton, Vitali Gretschko, Kemal Guler, Alok Gupta, Karla Hoffman, Maarten Janssen, Jayant Kalagnanam, Wolf Ketter, Ramayya Krishnan, Axel Ockenfels, Tuomas Sandholm, Achim Wambach, Christof Weinhardt, and Richard Steinberg for joint work and inspiring discussions. Lauren Cowles, Esther Miguéliz, and the team from Cambridge University Press provided outstanding support for the publication of the book.

I am particularly grateful to Paul Milgrom for all his kindness, advice, and guidance during my sabbatical at the Stanford Department of Economics. Special thanks also go to Dirk Bergemann, Bob Day, and Jacob Goeree for joint work, visits, and the many insights they offered. Last, and certainly not least, I thank my family for supporting me through this time-consuming project. I dedicate this book to them.