

Economic Modeling in Networking: A Primer

By Randall A. Berry and Ramesh Johari

Contents

1 Introduction	166
1.1 Endnotes	169
2 Welfare	171
2.1 Utility	172
2.2 Efficiency and Pareto Optimality	176
2.3 Fairness	178
2.4 Quasilinear Environments: The Role of Currency	184
2.5 Price Equilibria	187
2.6 Information	195
2.7 Endnotes	199
3 Equilibria	200
3.1 Static Games	201
3.2 Dominant Strategies	206
3.3 Nash Equilibria	209
3.4 Mixed Strategies	216
3.5 Bayesian Games	220
3.6 Repeated Games	227
3.7 Efficiency Loss	230
3.8 Endnotes	234

4 Externalities	236
4.1 Basic Definitions	237
4.2 Informational Externalities	241
4.3 Bargaining	245
4.4 Endnotes	248
5 Incentives	250
5.1 Regulation	251
5.2 Pigovian Taxes	254
5.3 The Vickrey-Clarke-Groves Mechanisms	260
5.4 Auctions	267
5.5 The Principal-Agent Problem	277
5.6 Endnotes	281
Acknowledgments	283
References	284

Economic Modeling in Networking: A Primer

Randall A. Berry¹ and Ramesh Johari²

¹ *Electrical Engineering and Computer Science, Northwestern University,
Evanston, IL 60208, United States, rberry@eecs.northwestern.edu*

² *Management Science and Engineering, Stanford University, Stanford,
CA 94305, United States, ramesh.johari@stanford.edu*

Abstract

In recent years, engineers have been increasingly called upon to have basic skills in economic modeling and game theory at their disposal for two related reasons. First, the economics of networks has a significant effect on the adoption and creation of network innovations, and second, and perhaps more importantly, engineered networks serve as the platform for many of our basic economic interactions today. This monograph aims to provide engineering students who have a basic training in economic modeling and game theory an understanding of where and when game theoretic models are employed, the assumptions underpinning key models, and conceptual insights that are broadly applicable.

1

Introduction

In recent years, network engineers have been increasingly called upon to have basic skills in economic modeling and game theory at their disposal. Economics chiefly concerns itself with the “production, distribution, and consumption of goods and services” (per Merriam–Webster’s Dictionary); game theory provides the theoretical foundations on which many economic models are built. Recent history has cast economics at the forefront of network engineering in particular in two ways. First, the economics of networks has a significant effect on the adoption and creation of network innovations (e.g., regulation of wireless spectrum — or the lack thereof — was a significant catalyst for the development of WiFi technologies). Second, and perhaps more importantly, engineered networks serve as the *platform* for many of our most basic economic interactions today. Despite this collision of economics and engineering, our own observation is that the typical networking student does not garner economic training until later in their graduate career — if at all.

Given the two facts in the previous paragraph, our position is that basic economic understanding is a critical tool in the arsenal of the modern student studying networks. Accordingly, this text is targeted

at providing a primer in economic modeling and game theory for such students in engineering, operations research, or computer science. Our focus is on providing a grounding for the basics, and particularly an understanding of where and when economic models are employed; the assumptions underpinning key models; and conceptual insights that are broadly applicable. We designed this monograph to be less of a mathematical or computational “how-to” guide to economic theory — there are innumerable textbooks available on the subject that address that requirement much more effectively. Rather, we felt a significant need to answer the following types of questions:

- (1) What is the difference between efficiency and fairness?
- (2) What assumptions lead to the result that markets are efficient?
- (3) When is a game theoretic model appropriate?
- (4) When is Nash equilibrium a reasonable solution concept?
- (5) Why is one auction format preferable to another?

More than just working with economic models, it is important for engineering students to be able to take a critical eye to such models as well. Our monograph is intended to help students develop that critical eye. In this sense, we view this monograph as a *complement* to a good introductory textbook or course on microeconomics or game theory; some references are provided in the endnotes to the introduction.

As a more specific example of the type of conceptual grounding we are after, we begin by noting that one can identify two distinct but interrelated approaches to the use of economic methods in the engineering sciences. First, we observe that in many instances economic models are used as a *semantic* device to aid in modeling an engineering problem. In this setting, the system under consideration is typically designed and implemented holistically; however, decentralized implementation requires that we understand the output of many interacting subcomponents. In some instances, viewing this interaction as a game can yield useful insight. In the second approach, game theoretic models are used as an *economic* device to explicitly capture incentives of a range of rational, self-interested parties. Here game theoretic models are used to predict the outcome of their interaction, and in some cases

design mechanisms to control that interaction. In contrast to the semantic approach, in the economic approach many aspects of the game are dictated by practical realities of the economic problem considered.

In the semantic approach, the usual goal is to leverage game theoretic techniques to design an optimal decentralized system. Congestion control is commonly cited as an example of this approach in networking. Here end-system controllers determine sending rates in response to congestion signals from the network; these signals can range from simple (packet loss due to buffer overflow) to complex (active queue management, or AQM, algorithms to determine explicit congestion notification, or ECN, marks). These algorithms can be naturally viewed in terms of a market, where the “buyers” are the end systems, and the “sellers” are the capacitated links of the network. Prices mediate supply and demand, and under certain conditions a market equilibrium exists and is efficient. Of course, the notion of “buying” and “selling” is purely semantic; however, it serves as a valuable (and powerful) guide that allows us to bring economic tools to bear on an engineering problem.

In the economic approach, by contrast, many features of the game are dictated by practical realities of the economic problem considered. A timely example is provided in recent studies of peer-to-peer (P2P) filesharing. Here there is a fundamental tension between the fact that a well-functioning filesharing system needs users to contribute content; but in the absence of any other incentive, users generally derive maximum benefit from downloading content, rather than contributing it. Note that in this setting the incentives are real: they are the incentives of the users contributing to the system, and any game theoretic model must start from this basic observation. At the same time, we emphasize that this approach can still involve an element of design: often we are interested in designing protocols or mechanisms that can provide the right incentives to self-interested agents. In that case, the economic approach shares much in common with the semantic approach to game theoretic modeling.

This kind of distinction may not be immediately evident to the network engineering student taking an economic modeling class for the first time. However, we believe it is a first order concern in formulating

any economic model in a networking setting. As we will see in the remainder of the text, it is a distinction that informs many of the concepts we present.

The remainder of the monograph is organized as follows. In Chapter 2, we introduce the basics of *welfare analysis* — the economic analysis of efficient allocation of resources among competing uses. In Chapter 3, we discuss game theoretic foundations, including a range of equilibrium concepts. In Chapter 4, we discuss *externalities* — the (positive or negative) impact of one individual’s actions on others. In Chapter 5, we discuss *mechanism design* — the design of economic environments (such as markets) that yield desirable outcomes.

A note on organization. Every section of every chapter is organized in an identical manner. Each section begins with a preamble setting the stage without technical details. The subsections then progress through *Methodology* (describing the technical approach); *Examples* (illustrating the methodology); and *Discussion* (describing key conceptual issues related to the section topic). The discussion subsections in particular play a significant role in our book: they delve beyond the “how” of economic modeling and game theory, and deal largely with the “why.” Each chapter concludes with a section of *Endnotes* that lists several references for the material discussed in the chapter. The literature on economics and game theory is vast. We have not tried to provide a comprehensive bibliography, but rather focus on a few key references to provide interested readers with a starting point for exploring the material in greater depth.

A note on terminology. Throughout the monograph, we refer to individuals as “players,” “agents,” or “users”; we use the terms interchangeably.

1.1 Endnotes

For more background on game theory we refer the readers to one of many books available on this subject such as Gibbons [12], Fudenberg and Tirole [11], Osborne and Rubinstein [32], Myerson [28], and Owen [33]. Also, Mas-Colell et al. [24] provides a good introduction to game theory as well many other aspects of microeconomics.

In theoretical computer science, algorithmic questions related to game theory have been increasingly studied; Nisan et al. [31] provides an excellent overview of this work. In communication networking, Walrand [42] is a good tutorial on economic issues relevant to communication networks, including situations where game theoretic modeling is used.

2

Welfare

In the context of resource allocation, *welfare analysis* uses an economic approach to study the overall benefit (or welfare) generated under alternative mechanisms for allocation of scarce resources. For our purposes, welfare analysis serves as a benchmark approach to resource allocation in engineered systems, and in particular allows us to introduce several essential concepts in economic modeling. We begin in Section 2.1 by introducing the notion of *utility*, the value that is derived by an individual from consumption of resources. Next, in Section 2.2, we discuss efficiency and define *Pareto optimality*, a way to measure the welfare of allocation choices. In Section 2.3, we discuss fairness considerations, and in particular how different notions of fairness lead to different ways of choosing between Pareto optimal outcomes. In Section 2.4, we introduce a particularly useful model of utility, where all market participants measure their utilities in a common unit of currency; these are known as *quasilinear environments*. In Section 2.5, we discuss the role of prices in mediating resource allocation through markets, and connect price equilibria to efficient market outcomes — as well as discuss some reasons why efficiency may fail to be achieved in markets. Finally, in Section 2.6, we briefly discuss the role of information in markets, and in particular discuss the welfare consequences of information asymmetries.

2.1 Utility

Utility provides a basic means of representing an individual's preferences for consumption of different allocations of goods or services. Utility is a central concept in the analysis of efficiency of economic systems, because it provides a metric with which to measure the satisfaction of an individual user. Welfare analysis is ultimately concerned with the utility levels generated by resource allocations across the system as a whole, but clearly a major building block in this exercise is to characterize the value of resources to a single user.

2.1.1 Methodology

Since utility functions are meant to capture an agent's preferences, we start by defining a *preference relation* for an agent. Suppose that there is a set J of resources available in the system, and let $\mathbf{x}_r = (x_{jr}, j \in J)$ denote a bundle of these resources allocated to user r . A preference relation, \succsim , is a binary relation used to model a user's preferences for different possible bundles, so that if $\mathbf{x}_r \succsim \mathbf{y}_r$, then user r finds \mathbf{x}_r at least as "good" as \mathbf{y}_r .

"Utility" refers to an assignment of value to each possible bundle that is aligned with an agent's preferences. Formally, a preference relation, \succsim , is represented by a *utility function*, $U_r(\mathbf{x}_r)$, if $U_r(\mathbf{x}_r) \geq U_r(\mathbf{y}_r)$ whenever $\mathbf{x}_r \succsim \mathbf{y}_r$, i.e., if a larger utility indicates that a bundle is more preferred by an agent. In some examples, the user's preferences and so their utility may depend only on a lower-dimensional function of the resource bundle; for example, in the case of resource allocation in networks, a user's utility may be a function only of the total rate allocation they receive. In these cases we adapt our notation accordingly.

There are two features of utility functions that are natural to consider in our context.

Monotonicity. We only consider utility functions that are nondecreasing, i.e., if every component of the allocated resource vector weakly increases, then the utility weakly increases as well. This means that every individual prefers more resources to less resources. A key implicit reason that monotonicity is plausible is the notion of *free disposal*: that excess resources can always be "disposed of" without any cost or

penalty. With free disposal, an individual is only ever better off with additional resources.

There are certain cases where consumption may be non-monotonic: for example, consumption of cigarettes may yield increased short-term utility at the expense of reduced long-term utility. In the case of online information, one reason for diminishing utility with increasing amounts of information may be our limited cognitive ability to search through it to find what we want (though this may be difficult to model). For our purposes, however, restriction to monotonic utility functions will suffice. Further, when utility depends only on a scalar quantity, we will often assume the utility function is strictly increasing for technical simplicity.

Concavity. Utility functions can take any shape or form, but one important distinction is between *concave* and *convex* utility functions. Concave utility functions represent diminishing marginal returns to increasing amounts of goods, while convex utility functions represent increasing marginal returns to increasing amounts of goods.

In the context of network resource allocation, there is an additional interpretation of concavity that is sometimes useful. Consider two users with the same utility function $U(x)$ as a function of a scalar data rate allocation x ; these users share a single link of unit capacity. We consider two different resource allocations. In the first, we randomly allocate the entire link to one of the two users. In the second, we split the link capacity equally between the two users. Which scenario do the users prefer?

Observe that if the users are making a phone call which requires the entire link for its minimum bit rate, then the second scenario is undesirable. On the other hand, if the users are each downloading a file, then the second scenario delivers a perfectly acceptable average rate (given that both users are contending for the link). The first (resp., second) case is often modeled by assuming that U is convex (resp., concave), because the expected utility to each user in the first allocation is *higher* (resp., lower) than the expected utility to each user in the second allocation. (Both results follow by Jensen's inequality.) For this reason, concave utility functions are sometimes thought to correspond to *elastic* traffic and applications, such as filesharing downloads; while

convex utility functions are often used to model applications with *inelastic* minimum data rate requirements, such as voice calls.

Finally, from a technical standpoint, we note that concavity plays an important role when we consider maximization of utility, given the technical simplicity of concave maximization problems.

2.1.2 Discussion

Preferences. If a utility represents an agent's preferences, a basic question is then: Do such representations always exist? The answer to this is "yes" provided that the user's preferences satisfy some additional assumptions. In particular, suppose an agent's preferences are complete (meaning a preference exists between any two bundles) and transitive (meaning that if $\mathbf{x}_r \succcurlyeq \mathbf{y}_r$ and $\mathbf{y}_r \succcurlyeq \mathbf{z}_r$, then $\mathbf{x}_r \succcurlyeq \mathbf{z}_r$); suppose also that the set of bundles is finite. Then it can be shown that such an agent's preferences can always be represented by a utility function. (For an infinite set of bundles, this is also true with an additional continuity assumption.)

A complete and transitive preference relation is also referred to as a *rational preference relation* and is assumed in nearly all economic modeling. Conversely, it can be shown that any utility function defined over a set of resource bundles gives a rational preference relation. Because of this fact, in many cases the assumption of rational preferences is implicit; often papers start by simply defining a utility function for each agent without mentioning the underlying preferences.

One might then ask: why are both concepts needed? Both serve a purpose. Preference relations are typically viewed as being more fundamental, and are something that can be at least partially inferred from the choices an agent makes; for example, we can repeatedly ask an agent which of two bundles she prefers. On the other hand, utility is a more compact representation and is more convenient for analytical work.

Rationality. A central theme in economic modeling is that the rational agent always acts to produce outcomes that are more preferred; in the context of utility functions, this means agents always act to maximize their own utility. Of course, this is a concept that is elegant in its simplicity, but routinely comes under fire for implausibility in a

range of settings. Perhaps the most common objection is that cognitive limitations often lead to suboptimal decisions.

There are two complementary responses to this issue. The first says that utilities are not “real”; rather, they are the modeler’s abstraction of how an agent behaves. In this view, what we are doing is taking data of observed behavior, and trying to model that behavior *as if* it came from optimization of an underlying utility function. This view is reasonable theoretically, but suffers in practical terms when utility representations that capture nuances of behavior become exceedingly complex.

The second response says that abstracted and simplified utility models are merely baselines to establish qualitative understanding of how incentives govern behavior. Often, when there are first order biases in behavior away from the fully rational model, these can be captured by considering heuristics of behavior that are perturbations to the initial utility model. For example, in dynamic environments, agents may act myopically rather than forecast the consequences of their actions in the future; this approach simplifies behavior yet still requires a model of utility that the agent maximizes myopically.

Comparing utility. When using a utility function to determine which of two bundles an agent prefers, only the *relative* ranking of the utilities matters, and not the *absolute* numeric values of the utility. In other words, whether agent r prefers \mathbf{x}_r to \mathbf{y}_r only depends on if $U_r(\mathbf{x}_r)$ is larger or smaller than $U_r(\mathbf{y}_r)$. For this reason such a utility function is referred to as an *ordinal utility*. Note that as a consequence, the ordinal utility representing an agent’s preferences is not unique. Indeed, if we set $V_r(\mathbf{x}_r) = f(U_r(\mathbf{x}_r))$ for any strictly increasing function f , then $V_r(\mathbf{x}_r)$ is another valid utility function that represents the same preferences as $U_r(\mathbf{x}_r)$.

Ordinal utility suffices for representing the preferences of one agent. However, one would often like to compare allocation decisions *across* agents. In such settings, it can be desirable for the units of utility to have some meaning. For example, even if two individuals both prefer two units of a good more than one unit of the same good, one may be interested in which agent values the one additional good the most, i.e., which agent receives the “most utility” from the good. Without additional assumptions, an ordinal utility does not allow us to answer such

questions, since the units of utility have no meaning. (This discussion focuses on modeling an agent's true economic preferences; when economic models are used as a semantic device, the modeler often designs utilities to reflect physical quantities and so they do have well defined units of measurement.) We discuss these issues further throughout the chapter; see in particular Section 2.4 for a particular setting where utility is measured in a common unit.

Risk aversion. The discussion of inelastic and elastic utility functions above is closely connected to the notion of *risk aversion* in microeconomic modeling. To see the connection, suppose we let $U(w)$ represent the utility to an individual of w units of currency (or wealth). Suppose we offer the individual a choice between (1) W units of wealth with certainty; or (2) a gamble that pays zero with probability $1/2$, or $2W$ with probability $1/2$. Note that the same kind of reasoning as above reveals that if U is concave, then the individual prefers the sure thing (W units of wealth with certainty); while if U is convex, the individual prefers the gamble. For this reason we say that an individual with concave utility function is *risk averse*, while an individual with a convex utility function is *risk seeking*. It is generally accepted that individuals tend to be risk averse at higher wealth levels, and risk neutral (or potentially even risk seeking) at lower wealth levels.

2.2 Efficiency and Pareto Optimality

A central question in any resource allocation problem is to define the notion of *efficiency*. To an engineer, efficiency typically means that all resources are fully utilized. However, this leaves many possibilities open; in particular, how are those resources allocated among agents in the system? It may be possible that some allocations that fully utilize resources lead to higher overall utility than others. In economics, efficiency is typically tied to the utility generated by resource allocations, through the notion of *Pareto optimality*.

2.2.1 Methodology

We require some notation and terminology. Let R be the set of users in a system, and let \mathcal{X} denote the set of all possible resource allocations.

Let U_r be the utility function of user r . Let $\mathbf{x} = (\mathbf{x}_r, r \in R)$ and $\mathbf{y} = (\mathbf{y}_r, r \in R)$ be two possible resource allocations to these users, i.e., two elements of \mathcal{X} . We say that \mathbf{x} *Pareto dominates* \mathbf{y} if $U_r(\mathbf{x}_r) \geq U_r(\mathbf{y}_r)$ for all r , and $U_s(\mathbf{x}_s) > U_s(\mathbf{y}_s)$ for at least one s in R . In other words, the allocation \mathbf{x} leaves everyone at least as well off (in utility terms), and at least one user strictly better off, than the allocation \mathbf{y} . From a system-wide (or “social”) standpoint, allocations that are Pareto dominated are undesirable, since improvements can be made to some users without affecting others.

This leads to the formal definition.

Definition 2.1. A resource allocation $\mathbf{x} \in \mathcal{X}$ is *Pareto optimal* if it is not Pareto dominated by any other allocation $\mathbf{y} \in \mathcal{X}$.

2.2.2 Examples

Example 2.1 (Single resource allocation). A key example we focus on throughout the manuscript is the basic problem of allocating a single resource among competing users. Formally, suppose that we are given a single resource of capacity C , and a set of users R , with user r having concave, strictly increasing utility function U_r as a function of the resource they are allocated. Assuming that the resource is *perfectly divisible*, i.e., the capacity can be divided arbitrarily among the agents, it follows that the set of possible resource allocations is given by

$$\mathcal{X} = \left\{ \mathbf{x} \geq 0 : \sum_r x_r \leq C \right\}.$$

It is clear that in this setting, the Pareto optimal allocations correspond exactly to the allocations where the entire link is utilized. Note, in particular, that there are many such allocations — and Pareto optimality by itself provides no guidance on how to choose between them.

Example 2.2 (Multiple resource allocation). Now suppose there are two resources, each with unit capacity, and two users. User 1 prefers

resource 1 to resource 2; while user 2 prefers resource 2 to resource 1. In particular, letting x_{jr} be the quantity of resource j allocated to user r , assume we have $U_1(x_{11}, x_{21}) = 2x_{11} + x_{21}$, while $U_2(x_{12}, x_{22}) = x_{12} + 2x_{22}$. Now consider two allocations that completely use all resources. In the first allocation \mathbf{x} , user 1 receives all of resource 1, and user 2 receives all of resource 2. In the second allocation \mathbf{y} , user 2 receives all of resource 1, and user 1 receives all of resource 2. The second allocation generates a utility of one unit to each user, while the first allocation generates a utility of two units to each user. Thus *both* allocations consume all resources; but only the *second* is Pareto optimal. We leave it as an exercise for the reader to determine whether there are any more Pareto optimal allocations in this setting.

2.2.3 Discussion

Economic efficiency vs. engineering efficiency. The examples highlight the important difference between economic efficiency and engineering efficiency. In the single resource setting, an economist's notion of efficiency corresponds with an engineer's notion of efficiency: the resource should be fully utilized. However, in general, economic efficiency is not only concerned with full utilization of available resources, but also optimal allocation of those resources among *competing* uses. This is the additional twist in the second example: it is essential to consider the value generated in considering which user receives each resource. Indeed, that challenge is at the heart of the economic approach to resource allocation, and this is why welfare analysis is so central in economics.

2.3 Fairness

As the examples in the previous section illustrate, there may be *many* Pareto optimal allocations. How do we choose among them? In this section, we will focus on *fairness* considerations in resource allocation. There are many competing notions of fairness; in this section, we focus on those notions that provide guidance for choosing between multiple possible Pareto optimal allocations. A central point of this section is to

illustrate the difference between Pareto optimality, which is an *objective* notion, and fairness, which is a *subjective* notion.

2.3.1 Methodology

We focus on a setting where the set of possible resource allocations in the system is given by \mathcal{X} ; we assume that \mathcal{X} is convex, compact, and nonempty. We assume that R users compete for resources, with user r 's utility given by $U_r(\mathbf{x}_r)$ at an allocation \mathbf{x}_r . For technical simplicity, we assume each user's utility is strictly concave and nondecreasing. (Strict concavity is not essential but simplifies the presentation below.)

We focus on a particular class of fairness criteria that are defined as follows. Let f be a strictly concave, strictly increasing function on the reals. In that case, it is straightforward to check that for each r , $f(U_r(\mathbf{x}_r))$ is a strictly concave function of \mathbf{x}_r . It follows that there exists a unique solution \mathbf{x}^* to the following optimization problem:

$$\begin{aligned} & \text{maximize} && \sum_r f(U_r(\mathbf{x}_r)) \\ & \text{subject to} && \mathbf{x} \in \mathcal{X}. \end{aligned}$$

Further, the solution \mathbf{x}^* is also Pareto optimal. (We leave verification of these facts as an exercise for the reader.)

Observe that *each* choice of f defines a potentially different choice of the Pareto optimal point \mathbf{x}^* . (Indeed, it can be shown that under our assumptions, *every* Pareto optimal point can be recovered through an appropriate choice of f .) We interpret f itself as encoding a fairness criterion in this way: the choice of f directly dictates the resulting “fair” allocation among all Pareto optimal points.

We note that, of course, there are many generalizations of this approach to fairness. For example, the function f may be user-dependent, or there may be a weight associated to each user in the objective function above. More generally, the aggregation function in the objective need not be a summation; we will see one such example below. In our presentation, for simplicity, we do not develop these generalizations.

Below we introduce three fairness criteria identified in this way (utilitarian, proportional, and α -fairness), and a fourth obtained as a limiting case (max-min fairness).

Utilitarian fairness. Perhaps the simplest choice of fairness is the case where f is the identity, in which case the Pareto optimal point that maximizes the *total* utility of the system is chosen. This is called the *utilitarian* notion of fairness; the name dates to the philosopher Jeremy Bentham, the founder of utilitarianism. An alternative interpretation is that utilitarian fairness implicitly assumes that all agents' utilities are measured in exactly the same units; and given this fact, the allocation should be chosen that maximizes the total "utility units" generated.

Proportional fairness. Utilitarian fairness can lead to allocations that remove users with lower utilities from consideration, in favor of those that generate higher utilities. One way of mitigating this effect is to *scale down* utility at higher resource allocations. This leads to the notion of *proportional* fairness, obtained by using $f(U) = \log U$.

The name "proportional fairness" is due to the following observation. Suppose for a moment that each user's utility depends only on a scalar quantity x_r , and that $U_r(x_r) = x_r$. Then if \mathbf{x}^* is the resulting optimal allocation, and \mathbf{x} is any other allocation, then the following inequality holds:

$$\sum_r \frac{x_r - x_r^*}{x_r^*} \leq 0.$$

In other words, the sum of proportional changes in the allocation across users cannot be positive. This condition can be verified from the optimality conditions for \mathbf{x}^* ; indeed, it is in fact an alternate characterization of the proportionally fair allocation. We leave the proof as an exercise for the reader.

α -fairness. Consider the following function f :

$$f(U) = \begin{cases} \frac{U^{1-\alpha}}{1-\alpha}, & \text{if } \alpha \neq 1; \\ \log U, & \text{if } \alpha = 1. \end{cases}$$

This utility function has the property that it is strictly concave and strictly increasing for all $\alpha \geq 0$. As α increases, the utility function exhibits progressively stronger decreasing marginal returns.

The resulting family of fairness notions, parameterized by α , is called *α -fairness*. This family is particularly useful because it includes

several special cases as a result. For example, when $\alpha = 0$, we recover utilitarian fairness; and when $\alpha = 1$, we recover proportional fairness. The case $\alpha = 2$ is sometimes called TCP fairness in the literature on congestion control in communication networks, because the allocation it leads to mimics the allocations obtained under the TCP congestion control protocol. Finally, an important special case is obtained when $\alpha \rightarrow \infty$; we discuss this next.

Max-min fairness. We conclude by discussing a fairness notion that does not fit in the framework above. In particular, \mathbf{x}^* is *max-min fair* if it solves the following optimization problem:

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \left[\min_r U_r(\mathbf{x}_r) \right].$$

It can be shown that this outcome is obtained as the limit of the α -fair allocation, as $\alpha \rightarrow \infty$.

Max-min fairness is sometimes called *Rawlsian fairness* after the philosopher John Rawls. It advocates for the protection of the utility of the least well off user in the system, regardless of whether a small utility change to this user might cause large changes in the utility of another user.

2.3.2 Example

To illustrate the differences in the fairness notions above, we consider the following example.

Example 2.3 (Network resource allocation). We consider a communication network resource allocation problem with three users and two resources, each of unit capacity. User 1 sends data only through resource 1. User 2 sends data only through resource 2. User 3 sends data through both resources 1 and 2. (See Figure 2.1.) Letting x_r be the rate allocation to user r , the two resources define two constraints on \mathbf{x} :

$$\mathcal{X} = \{\mathbf{x} \geq 0 : x_1 + x_3 \leq 1; x_2 + x_3 \leq 1\}.$$

Suppose the utility functions of all users are the identity utility function: $U_r(x_r) = x_r$ for all r .




Fig. 2.1 The network in Example 2.3.

It is then straightforward to check the following calculations:

- (1) The utilitarian allocation is $x_1 = x_2 = 1$, $x_3 = 0$.
- (2) The max-min fair allocation is $x_1 = x_2 = x_3 = 1/2$.
- (3) The proportionally fair allocation is $x_1 = x_2 = 2/3$, and $x_3 = 1/3$.

Thus the utilitarian allocation yields nothing to user 3, instead rewarding the shorter routes 1 and 2. The max-min fair allocation, by contrast, allocates the same data rate to user 3 as to users 1 and 2, despite the fact that user 3 uses twice as many resources. Proportional fairness falls in the middle: user 3 receives some data rate, but half as much as each of the other two users. In this sense, proportional fairness “interpolates” between the other two fairness notions; the same is true for other values of α with $0 < \alpha < \infty$.

2.3.3 Discussion

Subjective vs. objective: fairness vs. efficiency. As discussed at the beginning of the section, a central theme of our presentation is to clarify the stark contrast between efficiency (Pareto optimality) and fairness. Pareto optimality is an objective notion, that characterizes a set of points among which various fairness notions might choose. Fairness is inherently subjective, because it requires making a value judgment about *which* users will be preferred over others when resources are limited.

This is an important distinction. In particular, observe that there is no “tradeoff” between fairness and efficiency — the two are

complementary concepts. Different fairness notions correspond to different distributions of welfare across the population. Indeed, it is precisely because fairness is subjective that there can be significant debates on economic policy related to wealth redistribution, even among individuals that are equally well informed and qualified.

Fairness and protocol design. Fairness is perhaps one of the most difficult design decisions in resource allocation, because there need not be an objective criterion to guide the decision. An important misconception in designing resource allocation protocols is the notion that it is possible to *ignore* fairness in designing the protocol. Engineers may often feel that by following an established standard or norm, they are avoiding having to make a decision about the fairness criterion in allocating resources between competing uses; for example, every MAC protocol or congestion control protocol implicitly implements a fairness judgment across competing flows, applications, or users.

The key point is that *all allocation of scarce resources involves a value judgment*; the question is how this value judgment will be made. The protocol engineer must make sure they are aware of fairness criterion their deployment implements.

Axiomatic approaches to fairness. With so many different fairness notions, it seems natural to consider mathematical foundations for making choices among them. Two significant branches of literature consider axiomatic approaches to choosing among Pareto optimal allocations: the *social choice* literature, and the *axiomatic bargaining* literature. Both lines of literature aim to determine whether there are general mechanisms that choose outcomes consistent with particular fairness axioms. We particularly note an important result in the axiomatic bargaining literature relevant to our discussion here: the *Nash bargaining solution*. This solution was Nash's resolution of the problem of fair resource allocation; from a parsimonious set of axioms he derives a general rule for resource allocation. It can be shown that in the setting we discussed above, proportional fairness corresponds to a particular instantiation of the Nash bargaining solution. We emphasize that such axiomatic approaches do not remove the value judgment inherent in choosing a fairness criterion, but rather reframe this choice to one of selecting the set of fairness axioms.

Other approaches to fairness. Finally, we note that in our discussion, we focused on fairness notions that choose among Pareto optimal allocations. There are other fairness notions that may lead to allocations that are not even Pareto optimal. One example is *envy-freeness*. Formally, an allocation is envy-free if no user would strictly prefer to exchange their allocation with that given to another user. It is straightforward to construct examples of envy-free allocations that are not Pareto optimal; we leave this as an exercise for the reader.

2.4 Quasilinear Environments: The Role of Currency

To this point, we have treated utility as measured in arbitrary units. Among other things, this means we cannot directly compare the “value” that one user obtains from a quantity of resource with another user. This is one reason we resort to fairness notions in the first place: without any way to directly compare individual users, the set of possible Pareto optimal allocations can be quite large, necessitating a subjective choice among them. In this section, we show that the introduction of a common currency of measurement allows us to succinctly identify Pareto optimal allocations.

2.4.1 Methodology

In this section, we again assume that there are R users competing for resources, with the set of feasible allocations identified by \mathcal{X} . We assume that the system also consists of one additional good, a *currency*. We assume that users’ utilities are measured in units of currency. In particular, we assume that each user r ’s utility can be written as:

$$U_r(\mathbf{x}_r) + w_r,$$

where \mathbf{x}_r is the allocation of resources other than currency to user r , and w_r is her allocation of currency. Such utilities are called *quasilinear*, because they are linear in one of the goods (currency). Observe that in this definition, *every* user’s utility is evidently being measured in units of currency. When the utility is defined in this way, we refer to U_r as the *valuation* function for user r . Observe that U_r is itself measured in currency units as well.

Typically, for simplicity, we make two assumptions about the system when utilities are quasilinear. First, we assume that the initial endowment of currency to each individual is zero. Second, however, we also assume no credit constraints: we allow each individual to spend as much as they wish, subject to the constraint $\sum_r w_r \leq 0$. The last constraint ensures that an agent can only receive a *positive* amount of currency if the other agents spend at least that much in aggregate. If this inequality is strict, then some currency leaves the system. Note that, of course, if an individual spends currency, they incur a disutility given the structure of their utility function.

The possibility of monetary transfers greatly simplifies characterization of Pareto optimal allocations. The key intuition is that if an agent prefers a different allocation that makes another agent worse off, then a monetary transfer can be used as compensation. As long as the net combination of change in valuation and monetary transfer is positive for each agent, such changes in allocation will be Pareto improving. This intuition is formalized in the following result.

Theorem 2.1. Suppose all users have quasilinear utility. Then an allocation \mathbf{x}^* and vector of monetary transfers \mathbf{w}^* are Pareto optimal if and only if they satisfy the following two conditions:

$$\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathcal{X}} \sum_r U_r(\mathbf{x}_r);$$

$$\sum_r w_r^* = 0.$$

Proof. Suppose that \mathbf{x}^* and \mathbf{w}^* are Pareto optimal. If the second condition in the theorem does not hold, then currency leaves the system; thus in this case we can (feasibly) strictly improve the utility of any agent by returning this “lost” currency to them. Thus if \mathbf{x}^* and \mathbf{w}^* are Pareto optimal, then the second condition of the theorem must hold.

Now we focus on the first condition. Suppose that it does not hold; then there exists some other allocation \mathbf{x} with $\sum_r U_r(\mathbf{x}_r) > \sum_r U_r(\mathbf{x}_r^*)$. For each user r , define:

$$w'_r = U_r(\mathbf{x}_r^*) - U_r(\mathbf{x}_r).$$

Now consider the allocation \mathbf{x} together with the monetary transfers \mathbf{w}' . Observe that $\sum_r w'_r < 0$, and that the resulting overall utility to each user r is exactly $U_r(\mathbf{x}_r) + w'_r = U_r(\mathbf{x}_r^*)$. Thus the utility to each agent is unchanged, so all agents are exactly as well off as they were before; however, on balance, we have removed currency from the system. As in the previous paragraph, we can strictly improve the utility of any agent by returning the lost currency to them — violating Pareto optimality of \mathbf{x}^* and \mathbf{w}^* .

The reverse direction is similar, and left as an exercise for the reader. \square

The preceding theorem is an essential tool in the analysis of quasilinear environments. The basic point is that in quasilinear environments, *Pareto optimality corresponds to maximization of the sum of valuations* (analogous to what we called utilitarian fairness in the preceding section). Note that this somewhat eases fairness considerations, particularly in contexts where the aggregate valuation maximizing resource allocation is uniquely defined.

Often, in an abuse of terminology, it is simply said that an allocation \mathbf{x}^* is Pareto optimal if it maximizes the sum of valuations; in other words, any vector \mathbf{w}^* of monetary transfers that sums to zero can be chosen, and the resulting pair \mathbf{x}^* and \mathbf{w}^* would be Pareto optimal in the formal sense described above.

An additional value of the preceding result is that it gives us a single metric by which to measure the economic efficiency of the system. In general, entire utility vectors must be compared to determine Pareto optimality. By contrast, in quasilinear environments, we can measure the welfare of the system through the single metric of the sum of valuations. Indeed, the sum of valuations is sometimes directly referred to as the *aggregate welfare* for quasilinear environments. Note that this is the explanation of the term “welfare” in the title of this chapter.

2.4.2 Discussion

Uniqueness of optimal solutions. Quasilinear environments are particularly easy to work with when the aggregate utility optimization problem is strictly concave, since in this case the Pareto optimal allocation is

unique. In more complex settings, it may be difficult to solve for the Pareto optimal allocations. One important example is provided by combinatorial auctions; these are auctions where bidders compete for bundles of goods, and the valuation function of a bidder can depend in complex ways on the exact composition of the bundle received. The combinatorial structure in bidders' valuations can make the resulting optimization problem quite complex.

First-best optimality. We note here that maximization of the sum of valuations provides a natural benchmark against which to compare *any* scheme for resource allocation. For this reason the Pareto optimal utility (and resulting welfare) are often referred to as the “first-best” outcome, as compared to “second-best” outcomes arising from alternative mechanisms. Depending on the context, second-best solutions typically suffer some efficiency loss relative to the first-best outcome, but may have other properties that make them more practically appealing.

Plausibility of quasilinearity. Even in the presence of a currency, there are reasonable questions regarding whether quasilinearity is plausible. In some cases, it is a matter of convenience to focus on quasilinear utilities. This is certainly the case, for example, in most models in auction theory; see Chapter 5. Generally, such an assumption is unreasonable when there are significant differences in the marginal utility of wealth to an agent at different wealth levels. For example, at high wealth levels, the marginal effect of an additional unit of currency is generally lower than at low wealth levels. An even more basic objection can be raised in the context of network protocol design, where no obvious currency or direct mechanism of comparison exists between different users or different applications. In such settings, particularly when economic models are only used as a semantic tool, it may be reasonable to assume preferences that are quasilinear in some other variable (such as delay or throughput). Ultimately, any time that quasilinearity becomes an unreasonable assumption, the door is reopened to a discussion of suitable fairness criteria.

2.5 Price Equilibria

In the preceding sections, we have discussed the notion of efficiency for an economy; notably, in the case of quasilinear economies, we

discovered that Pareto optimality corresponds to maximization of aggregate welfare. However, our discussion thus far has not addressed a significant question: *how do individuals find Pareto optimal allocations?* In this section, we investigate how *prices* can play the role of mediating variables that help lead a system to Pareto optimal outcomes, under the right conditions.

2.5.1 Methodology

For simplicity in this section we consider a quasilinear environment, i.e., all participants in the system have their utilities measured in monetary units. We consider a slightly richer model in this section than in preceding sections. We continue to consider a setting where R users compete for resources; let J denote the number of distinct resources. As before, let x_{jr} denote the amount of good j consumed by consumer r . Due to quasilinearity, if consumer r receives a bundle of goods $\mathbf{x}_r = (x_{jr}, j = 1, \dots, J)$, and receives w_r in currency, then user r 's utility is:

$$U_r(\mathbf{x}_r) + w_r. \quad (2.1)$$

Throughout, we assume that U_r is concave, nondecreasing in \mathbf{x}_r , continuous, and continuously differentiable.

Up to this point the system looks identical to that in the preceding sections. However, in this section, we also add *producers* to the mix. Producers are able to create resources for consumption by users. Formally, we assume there are S producers and let y_{js} denote the amount of resource j produced by firm s . If producer s receives revenue t_s , and produces a vector $\mathbf{y}_s = (y_{js}, j = 1, \dots, J)$, then we assume producer s earns a profit given by:

$$t_s - C_s(\mathbf{y}_s). \quad (2.2)$$

Here C_s is the *production cost* function of producer s . Throughout, we assume that C_s is convex, continuous, and continuously differentiable. In addition, for technical simplicity we assume that $C_s(\mathbf{y}_s)$ is strictly increasing in each coordinate y_{js} (though this is not essential).

Suppose that the combined system of users and producers can destroy any created resources (including money) at zero cost, but that no goods (including money) can be injected into the economy. Then by arguing as in the preceding section, we can conclude that the allocation (\mathbf{x}, \mathbf{y}) together with transfers (\mathbf{m}, \mathbf{t}) constitute a Pareto optimal operating point if and only if:

(1) The pair (\mathbf{x}, \mathbf{y}) solves the following optimization problem:

$$\text{maximize } \sum_r U_r(\mathbf{x}_r) - \sum_s C_s(\mathbf{y}_s) \quad (2.3)$$

$$\text{subject to } \sum_r \mathbf{x}_r = \sum_s \mathbf{y}_s, \mathbf{x}, \mathbf{y} \geq 0. \quad (2.4)$$

(2) The transfers balance: $\sum_r m_r = \sum_s t_s$.

As discussed in the preceding section, we sometimes abuse terminology by saying that (\mathbf{x}, \mathbf{y}) are *efficient* if they solve the optimization problem (2.3)–(2.4) (since in this case monetary transfers that balance would yield Pareto optimality in the usual sense).

Let $W(\mathbf{x}, \mathbf{y}) = \sum_r U_r(\mathbf{x}_r) - \sum_s C_s(\mathbf{y}_s)$. In addition to the relevant concavity/convexity assumptions above, we also assume now that $W(\mathbf{x}, \mathbf{y})$ is *coercive*, i.e., we assume that as $\|\mathbf{y}\| \rightarrow \infty$, $W(\mathbf{x}, \mathbf{y}) \rightarrow -\infty$ as well. Under these assumptions, there exists at least one optimal solution to (2.3)–(2.4).

Let's now consider how the economy of users and producers might find such an operating point. We consider an abstraction where a price vector mediates exchange between users and producers, i.e., prices determine the amount of currency per unit resource that users transfer to the producers. We aim to describe a “stable point” of such a market. In particular, we expect two conditions to hold: first, that given prices, both users and producers have optimized their net utility or profit; and second, that given their decisions, the use of resources matches the production of resources (i.e., “the market clears”).

Formally, we have the following definition of a *competitive equilibrium*.

Definition 2.2. A pair of vectors $\mathbf{x}^e = (\mathbf{x}_r^e, r = 1, \dots, N)$, and $\mathbf{y}^e = (\mathbf{y}_s^e, s = 1, \dots, M)$, together with prices $\mathbf{p}^e = (p_j, j = 1, \dots, J)$ constitute a *competitive equilibrium* if:

- (1) All consumers have optimized utility given prices: for all r , there holds:

$$\mathbf{x}_r^e \in \arg \max_{\mathbf{x}_r \geq 0} U_r(\mathbf{x}_r) - \mathbf{p}^\top \mathbf{x}_r.$$

- (2) All producers have maximized profit given prices: for all s , there holds:

$$\mathbf{y}_s^e \in \arg \max_{\mathbf{y}_s \geq 0} \mathbf{p}^\top \mathbf{y}_s - C_s(\mathbf{y}_s).$$

- (3) The market clears: for all j there holds:

$$\sum_r x_{jr}^e = \sum_s y_{js}^e.$$

A critical feature of the definition of competitive equilibrium is that all players are *price taking*. In other words, they optimize holding the prices fixed, and do not forecast that their actions will have any effect on the prices in the market. We return to this assumption in the discussion below.

The following theorem summarizes some key properties of competitive equilibria.

Theorem 2.2. Under the assumptions above, there exists at least one competitive equilibrium $(\mathbf{x}^e, \mathbf{y}^e, \mathbf{p}^e)$. Further, at any competitive equilibrium the resulting allocation $(\mathbf{x}^e, \mathbf{y}^e)$ is efficient. Conversely, given any efficient allocation $(\mathbf{x}^*, \mathbf{y}^*)$, there exists a vector of prices \mathbf{p}^* such that $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{p}^*)$ is a competitive equilibrium.

Proof. Our proof proceeds by identifying competitive equilibria with solutions to the problem (2.3)–(2.4).

Let $\boldsymbol{\lambda}$ be a vector of Lagrange multipliers for the constraint (2.4); the Lagrangian for (2.3)–(2.4) is:

$$\mathcal{L} = W(\mathbf{x}, \mathbf{y}) - \boldsymbol{\lambda}^T \left(\sum_r \mathbf{x}_r - \sum_s \mathbf{y}_s \right).$$

Under our concavity/convexity assumptions, first-order conditions are necessary and sufficient for optimality. We conclude that $(\mathbf{x}^*, \mathbf{y}^*)$ is efficient if and only if there exists a Lagrange multiplier vector $\boldsymbol{\lambda}$ such that:

$$\text{For all } r : \frac{\partial U_r}{\partial x_{jr}}(\mathbf{x}_r^*) = \lambda_j, \quad \text{if } x_{jr}^* > 0, \quad (2.5)$$

$$\leq \lambda_j, \quad \text{if } x_{jr}^* = 0; \quad (2.6)$$

$$\text{For all } s : \frac{\partial C_s}{\partial y_{js}}(\mathbf{y}_s^*) = \lambda_j, \quad \text{if } y_{js}^* > 0; \quad (2.7)$$

$$\leq \lambda_j, \quad \text{if } y_{js}^* = 0; \quad (2.8)$$

$$\sum_r \mathbf{x}_r^* = \sum_s \mathbf{y}_s^*. \quad (2.9)$$

Note that (2.5)–(2.6) ensure that each consumer r has maximized utility given the prices $\boldsymbol{\lambda}$; that (2.7)–(2.8) ensure that each firm has maximized profit given the prices $\boldsymbol{\lambda}$; and that (2.9) ensures that the market clears. In other words, if (\mathbf{x}, \mathbf{y}) is an efficient allocation, then there exist prices $\boldsymbol{\lambda}$ such that $(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda})$ is a competitive equilibrium. Conversely, if $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{p}^*)$ is a competitive equilibrium, the same argument establishes that $(\mathbf{x}^*, \mathbf{y}^*)$ must be efficient. Existence follows since there exists at least one efficient allocation (as W is coercive). \square

The preceding theorem reveals two important facts about competitive equilibria: first, that any competitive equilibrium is efficient; and second, that given any efficient allocation, there exist prices that “support” this allocation as a competitive equilibrium outcome. The first of these observations is a special case of *the first fundamental theorem of welfare economics*, and the second observation is a special case of *the second fundamental theorem of welfare economics*. Taken together, the welfare theorems are often a justification for why we might expect

price equilibria in markets to yield efficiency — though it is important to note that there are many potential sources of inefficiency in markets, as we will discuss throughout the rest of the monograph.

2.5.2 Example

We conclude with an example demonstrating how price equilibria might be interpreted in the context of resource allocation in networks.

Example 2.4 (Network resource allocation). Suppose that R users want to send data across a network consisting of J resources (i.e., J links). We uniquely identify a user r with a subset of the resources, constituting a path. Further, we assume that user r receives a (monetary) value $U_r(x_r)$ from sending data at rate x_r along her path. We use the notation “ $j \in r$ ” to denote the fact that resource j lies on the path of user r . Let $f_j = \sum_{r:j \in r} x_r$ denote the total rate on link j . We assume that link j has a *delay function*, $\ell_j(f_j)$, that denotes the delay per unit rate felt by each user sending traffic through link j . Our only assumption is that delay is measured in monetary units, i.e., one unit of delay equals one monetary unit of utility lost.

Let us view this system as an economy with quasilinear preferences, by treating each link as a separate producer j , “producing” its own resource — capacity at link j . The key step we take is that we *define* the cost function of link j as $C_j(f_j) = f_j \ell_j(f_j)$. Note that this is the total delay experienced by flows at resource j . Throughout this section, we assume that ℓ_j is convex and strictly increasing, $U_r'(0) = \infty$, and U_r is strictly concave and strictly increasing.

Under our assumptions, a competitive equilibrium $(\mathbf{x}, \mathbf{f}, \boldsymbol{\mu})$ exists, and satisfies the following three conditions:

$$\begin{aligned} U_r'(x_r) &= \sum_{j \in r} \mu_j; \\ C_j'(f_j) &= \ell_j(f_j) + f_j \ell_j'(f_j) = \mu_j; \\ \sum_{r:j \in r} x_r &= f_j. \end{aligned}$$

Further, the resulting rate allocation solves the following optimization problem:

$$\begin{aligned} & \text{maximize} && \sum_r U_r(x_r) - \sum_j f_j \ell_j(f_j) \\ & \text{subject to} && \sum_{r:j \in r} x_r = f_j, \quad \text{for all } j; \quad \mathbf{x}, \mathbf{f} \geq 0. \end{aligned}$$

Now there are several ways to leverage this formulation. By far the most common is *reverse engineering*: we can exploit this price-based interpretation of network resource allocation as a way to understand what network protocols are doing. This approach typically involves viewing network protocols as “optimizing” utility for a flow, given congestion signals (e.g., delay) received from the network.

We conclude the example by observing that the prices, $\boldsymbol{\mu}$, in the resulting competitive equilibrium are *not* just the delays at each link. In particular, the price at link j is the delay $\ell_j(f_j)$ plus an additional term, $f_j \ell'_j(f_j)$. This latter term is sometimes called a *Pigovian tax*, after the economist Pigou; it is charged to account for the fact that additional flow sent by user r through link j not only increases the delay to user r , *but also to all other users of link j as well*. We say that user r imposes a *negative externality* on the other users at link j . The Pigovian tax ensures that user r “internalizes” this externality when she determines her sending rate. (See Chapter 4, and Section 5.2 in particular, for more detail on externalities and Pigovian taxes.)

2.5.3 Discussion

Price taking. As noted above, a critical assumption in the competitive equilibrium definition is that users and producers are *price takers*, i.e., that they hold prices constant while they optimize. Generally speaking, such an assumption is reasonable only when users and producers are “small” relative to the entire market, so that their individual actions are unlikely to affect prices in the market as a whole. As an example at the other extreme, suppose that there were only *one* producer in the entire marketplace. Then informally, we should expect this producer to be able to *set* its own price for the rest of the market, whether directly or

indirectly (through a choice of production quantity). This is an example of “market power,” where one or a few participants are large enough that their actions can directly influence prices in the market place. In such a setting, the price taking assumption underlying competitive equilibrium breaks down. Most importantly, models that properly account for strategic behavior in the presence of market power typically predict a loss of efficiency due to the price-anticipating nature of dominant users and/or producers. (See Chapter 3 for a more detailed discussion of such models.) This is an important source of deviations away from the predictions of the welfare theorems.

Convergence to competitive equilibrium. Competitive equilibrium is a *static* concept, and does not describe the process by which markets form prices over time, i.e., the *dynamics* of the market. There are several different lenses through which to view market dynamics. In this remark we briefly discuss two: a computational view, and a strategic view.

We first consider whether competitive equilibria are easily computable. This question has proven to be of significant interest in network resource allocation, where decentralized computation of competitive equilibria has been used to model and provide design guidance for rate control protocols in networks. There are at least two processes here that are often studied: *quantity adjustment processes* (sometimes called “primal” algorithms) and *price adjustment processes* (also called “dual” algorithms, or tâtonnement processes). Under quantity adjustment, end users adjust their flow rates, and resources feed back congestion signals to users as a function of current flows; users then adjust rates up or down until their marginal utility matches the total route price. Under price adjustment, end users instantaneously choose rates so that their marginal utility matches the currently observed route price; the resources then adjust prices up or down to bring marginal cost at the resulting flow rates in line with prices. Both algorithms can be shown to converge to competitive equilibrium, under appropriate assumptions.

Notice, however, that this dynamic view of markets is somewhat limited. In particular, it assumes that users and producers react in a myopic way to current price signals from the market. This view is potentially useful in designing algorithms (such as congestion control

protocols); but is less useful in describing actual market dynamics with strategic actors. In the latter setting, we must also account for the complex ways in which individuals might reason about *future* market dynamics given the present state of the market. Studying such models of market dynamics remains an active area of investigation, and a detailed discussion is outside the scope of this text. However, we emphasize that such dynamics provide yet another example of how real markets can yield deviations away from the simple, static notion of equilibrium and efficiency discussed above.

Pigovian taxes: Inefficiency and Braess' paradox. We conclude the section with a further discussion of Pigovian taxes. As noted above, in the network resource allocation example, prices are set so that users optimize taking into account the harm they inflict on others using the same resources. It is also worth considering how the market might behave if prices were more naively set to *just* the delays at each link, excluding the Pigovian tax. In this case it is possible to show (using an argument similar to Theorem 2.2) that a corresponding “equilibrium” still exists; however, since it is an equilibrium with the “wrong” prices, the resulting allocation will not be efficient in general. Extensive recent work has been devoted to studying the extent of this inefficiency.

Even more strikingly, this allocation can behave in unusual ways as we alter the underlying network. For instance, a famous example known as *Braess' paradox* shows that if prices are set to delays (and do not include the Pigovian tax) then *adding* a link to the network can actually *worsen* the utility for all users! Although surprising on the surface, the result begins to be more plausible when one observes that excluding the Pigovian tax does not ensure that efficient allocations are reached. When prices exclude the Pigovian tax, the resulting allocations can be inefficient; and Braess' paradox shows the additional insight that such allocations need not even become more efficient when links are added to the network.

2.6 Information

In all the settings discussed thus far, agents share the same information about the system they interact with. In particular, in the model of price

equilibria in the preceding section, agents did not face uncertainty in any part of their decision process: agents had known utility or cost, and optimized given a known price. In this section, we show that a lack of complete information can lead to potential inefficiencies in market outcomes, depending on the mechanism of trade that is used.

2.6.1 Methodology

We consider a setting motivated by trade in online marketplaces such as eBay. For simplicity, suppose that one buyer wants to purchase a single unit of a good (e.g., a rare coin). Suppose this buyer can purchase the good from a seller, who has one unit of the item for sale. We make two further assumptions. First, we assume that the good has a value of v to the seller; in the case of rare coins, for example, this might be viewed as the intrinsic book value of the coin. Second, we assume that the buyer values the good somewhat more than the seller, but prefers goods with higher intrinsic value. We model this by assuming the value to the buyer of the seller's good is αv , where $\alpha > 1$.

Let's start by considering whether the economy consisting of the buyer and seller has a competitive equilibrium. Since $\alpha > 1$, the social welfare maximizing (and efficient) allocation is to transfer the good to the buyer: this generates welfare αv , whereas letting the seller keep the good only generates welfare v . Such an allocation is "supported" by any price p such that $v < p < \alpha v$, since at such a price the seller will choose to sell the good, and the buyer will choose to purchase the good. Thus any such price together with the allocation of the good to the buyer constitutes a competitive equilibrium and an efficient outcome.

We now change the model of the preceding paragraph. In particular, suppose that the seller's value v is a uniform random variable on $[0, 1]$ — observed by the seller, but unknown to the buyer. The motivation for such a model is clear in markets like eBay: since the seller holds the good, she typically has a much better understanding of the value of the good than the buyer. For simplicity, we continue to imagine that the user values the good at a multiple of v , i.e., the value to the buyer is αv where $\alpha > 1$. We emphasize that with such a model, the value is *unknown* to the buyer; if v were known to the buyer, *the*

efficient outcome is exactly the same as before (since $\alpha > 1$): the good is transferred to the buyer from the seller.

We now develop a notion of price equilibrium for this model. In particular, suppose that a price p is set in the market. In this case, the producer will sell the good if $p > v$; not sell if $p < v$; and be indifferent between selling or not if $p = v$. It seems reasonable that *the buyer should incorporate this information in making a purchasing decision*; this is typically known as “rational expectations.” In particular, the buyer should infer that if a purchase can be made a price p , the *expected* value to the seller must be $E[v|v < p] = p/2$, and so the expected value to the buyer is $\alpha p/2$. This analysis suggests that if the buyer maximizes net expected payoff, she will buy if and only if $\alpha p/2 - p > 0$, i.e., if and only if $\alpha > 2$. In particular, we reach the following conclusion: *when $1 < \alpha < 2$, the buyer will never choose to purchase at any price!*

This phenomenon is known as *market failure*; the particular example described here is referred to as a “market for lemons,” referring to the fact that the buyer infers that the goods sold at any given price are “lemons” — those with low expected values relative to the price. Despite the fact that it is known in advance that the efficient outcome involves trade from the seller to the buyer, there does not exist any price where the buyer would be willing to purchase from the seller. The buyer reasons that at any given price, the seller would sell only in those instances where v is low — and therefore the value to the buyer is also low. In other words, the uncertainty in the value to the buyer combines with her (correct) forecast that the good will only be sold when its value is low to lead to market failure. We emphasize that the main goal of this exercise is to illustrate how an information asymmetry (in this case, between seller and buyer) can lead to stark inefficiencies in market outcomes relative to the competitive equilibrium benchmark of the preceding section.

2.6.2 Discussion

Contracting and completeness of markets. Observe that one reason why there is market failure in the setting above is that the market structure is quite rigid. Consider an alternate market procedure, where the buyer

and seller *ex ante* (before the value of the good is revealed) agree to a price for each possible value that might arise; such a process yields a *price functional* $p(v)$, as a function of the value of the good that is sold. (Note that such agreement requires that the buyer be able to validate the value of the good; for example, the buyer and seller might agree to an independent appraisal of value.) In this case, it is clear that with any price functional where $v < p(v) < \alpha v$ for all $v \in [0, 1]$, the Pareto optimal outcome is achieved — the seller would choose to sell, and the buyer would choose to buy.

Why is the price functional different from the single price chosen above? The reason is that the price functional allows the buyer and seller to contract over every possible *contingency* — i.e., every possible value of v that might arise. By contrast, in the market for lemons, the buyer must commit to buy or sell without any assurance of what value of v might arise. In settings with asymmetric information, when contracting is possible over every contingency, markets are said to be *complete*. This example illustrates a general principle that completeness of markets recovers Pareto optimality of price equilibria, even with asymmetric information. However, the example also illustrates why completeness is rarely achieved in practice: such a market has much higher overhead than a simpler market where a single price mediates trade (rather than a price functional). In other words, in reality markets are typically *incomplete*.

Adverse selection. The market for lemons exhibits a common feature of markets with asymmetric information. In particular, consider the *seller's* decision problem at a given price level. Observe that at any given price level, it is precisely the *lowest* quality goods that would be offered for sale by the seller (since this is the only case where it is worth it for the seller to sell). In other words, the goods that are offered for sale are *adversely* selected from the perspective of the buyer. *Adverse selection* is a term generally used to refer to the fact that when individuals make decisions in markets without full information at their disposal, they may feel regret *ex post* when the truth is revealed to them. Since economic models of behavior postulate that individuals rationally factor in this potential regret, adverse selection risk typically leads market participants to be more cautious than they would be in

the presence of full information. In the example in this section, it is this cautiousness caused by adverse selection risk that leads the buyer to decline trade entirely when $\alpha < 2$.

2.7 Endnotes

Many of the topics in this chapter are covered in greater detail in standard textbooks on microeconomic theory; see, e.g., Mas-Colell et al. [24] or Varian [39]. Some specific references related to material in each section follow.

Utility. The notion of elastic and inelastic traffic was introduced by Shenker [38].

Fairness. In the networking literature, the notion of proportional fairness was introduced by Kelly [18], while the notion of α -fairness was introduced by Mo and Walrand [26]. Max-min fairness has a longer history of use in the networking literature; see, e.g., Chapter 6 of Bertsekas and Gallager [5].

Price equilibria. The reader should consult Chapters 15–18 of the textbook by Mas-Colell et al. [24] for more details on price equilibria, including price equilibria in environments that are not quasilinear (known as *general equilibrium*).

In the networking literature, Kelly [18] and Kelly et al. [19] initiated the study of network resource allocation via price equilibria; the second of these papers focuses on dynamic processes for convergence to equilibrium. See also the work of Low and Lapsley [22] and the survey by Shakkottai and Srikant [37].

See Roughgarden and Tardos [36] for a discussion of inefficiency in routing problems; the interested reader may also wish to see Pigou's early work [35].

Information. The market for lemons presented in the text is a slight variation of the work in Akerlof's celebrated paper [1]. See also Chapter 19 in the textbook by Mas-Colell et al. [24] for further discussion on price equilibria in markets with incomplete information.

3

Equilibria

In the previous chapter we argued that in many cases a fair and efficient allocation of resources can be obtained when agents individually maximize their own net benefit given an appropriate set of resource prices. A key assumption in this story is that agents are *price takers*, i.e., they do not account for any effect their actions have on an underlying price setting process. Price taking agents can then be viewed as if each is solving an optimization problem in isolation.

In some settings, this may be a questionable assumption. If an agent knows that its behavior will change the price, then why should it not account for this effect? Relaxing the price taking assumption requires an agent to consider not only her own action, but the actions of all other agents, since prices are determined by the joint behavior of all agents. The corresponding optimization problems faced by each agent then become interdependent. Modeling such situations falls into the purview of *game theory*, which we begin discussing in this chapter.

Game theory provides a well established framework for modeling the interactions of multiple decision makers. Indeed, as we indicated in Chapter 1, this framework underlies much of modern economic theory. We begin in Section 3.1 by introducing a basic model for a *static game*. Next, we turn to discussing various approaches for “solving” a

game, beginning with the concept of *dominant strategies* in Section 3.2. We introduce the celebrated notion of *Nash equilibrium* in Section 3.3. Our initial discussion focuses on games in which agents choose a single deterministic action; in Section 3.4, we discuss the role of randomization in games and *mixed strategy* equilibria. In Section 3.6, we discuss dynamics in games through the model of *repeated games*. Finally, in Section 3.5, we highlight the role of information, and discuss the framework of *Bayesian games* as a way of modeling games with incomplete information.

3.1 Static Games

In game theory a “game” refers to any situation in which multiple agents or players interact. A model of a game specifies *how* they interact (e.g., what are their possible actions, when can they act, etc.), and also specifies how they value the outcomes of this interaction. In this section we discuss a basic class of such models, called *static* or *one-shot* games, in which the players interact once and simultaneously make a decision. Though this is clearly an over simplified model of many situations, it can still provide much insight into how agents may interact, and forms the basis for many generalizations that allow for richer interactions.

3.1.1 Methodology

Formally, a static game is defined as follows¹:

Definition 3.1. A *static game* G is a 3-tuple $(R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ where:

- (1) R is a finite set of players;
 - (2) for all $r \in R$, S_r is the *strategy set* of player r ; and
 - (3) for all $r \in R$, $\Pi_r : \prod_{s=1}^R S_s \mapsto \mathbb{R}$ is the *payoff function* of user r .
-

¹More precisely, this is a *strategic* or *normal form* description of a game. An alternative description, called the *extensive form* is useful to model more general classes of games, and in particular to capture dynamic behavior in a game.

In this definition, each player or agent $r \in R$ is an individual decision maker who must choose an *action* s_r from a set of possible actions given by their strategy set S_r . An outcome \mathbf{s} of the game is a choice of action for each agent $r \in R$. The payoff function Π_r models agent r 's preferences over outcomes and can be viewed as a utility function as discussed in the previous section. Agent r 's goal is to maximize her payoff $\Pi_r(\mathbf{s}) = \Pi(s_r, \mathbf{s}_{-r})$, where \mathbf{s}_{-r} denotes the set of actions of all players other than r . This notation emphasizes the fact that a player's payoff depends on both her own action (s_r) and the actions of all of the other players (\mathbf{s}_{-r}). We also denote the set of all possible actions for these players by S_{-r} , and let $S = \prod_{s=1}^R S_s$ denote the set of all possible *joint* action profiles across all players. We sometimes also refer to a particular choice of $\mathbf{s} \in S$ as an *outcome* of the game.

3.1.2 Examples

We next give two examples of static games.

Example 3.1 (Peering). Consider an interconnection, or *peering*, game between two Internet service providers (ISPs). Suppose that two providers 1 and 2 have agreed to connect to each other in New York and San Francisco (see Figure 3.1). Suppose that 1 MB of traffic is originating with a customer of ISP 1 in Chicago, destined for a customer of ISP 2 in San Francisco; and similarly, 1 MB of traffic is originating

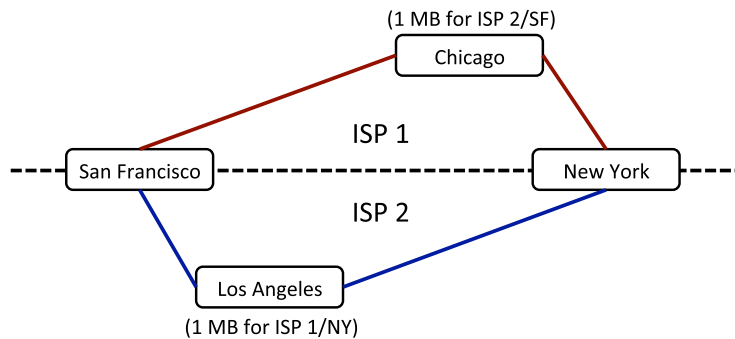


Fig. 3.1 The peering game in Example 3.1.

with a customer of ISP 2 in Los Angeles, destined for a customer of ISP 1 in New York.

We assume that routing costs to the ISPs are proportional to distance traveled; in particular, assume each unit of traffic that traverses a short link (resp., long link) incurs a cost of 1 unit (resp., 2 units). Each ISP can choose to send traffic to the other ISP across one of two *peering points*: San Francisco and New York. Note that New York is the nearer exit for ISP 1, and San Francisco is the nearer exit for ISP 2; we use “near” and “far” to denote the actions of routing to the nearest and furthest peering points, respectively. We can model this as a static game in which the two ISPs are the players ($R = \{1, 2\}$), each with a strategy set given by $S_r = \{\text{near}, \text{far}\}$. The payoffs for each agent can be represented as the follows:

		ISP 2	
		near	far
ISP 1	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

In this matrix, (x, y) means that the payoff x is obtained by ISP 1, and the payoff y is obtained by ISP 2, when the corresponding action pair is played. Thus, for example, when ISP 1 routes to the nearest exit, while ISP 2 routes to the furthest exit, then ISP 1 incurs a cost of only 1 unit, while ISP 2 incurs a cost of 2 units (to route its own traffic to New York) plus 3 units (to reroute the incoming traffic from ISP 1 from New York to San Francisco).

Example 3.2 (Allocation of a congestible resource). In the peering game each player’s strategy space is finite and so we can represent the payoffs of each player using a matrix. Now we give an example of a game in which the strategy spaces are infinite sets. This example involves sharing a single resource among R users. In this game, each

user is a player and their action is the quantity of resource they consume x_r (e.g., their rate). To model the payoffs, as in Example 2.4, we assume that each user derives a utility $U_r(x_r)$ from the resource it is allocated, and that the resource is *congestible* so that as the total usage increases, each user experiences a disutility per unit resource given by $\ell(y)$, where ℓ is a convex and strictly increasing function of the total allocation $y = \sum_r x_r$. With these assumptions, the payoff to user r is:

$$\Pi_r(x_r, \mathbf{x}_{-r}) = U_r(x_r) - x_r \ell \left(\sum_s x_s \right). \quad (3.1)$$

Note that a user's payoff depends implicitly on the transmission rate decisions of all other players, through the common congestion function.

If we assume that players can choose any nonnegative allocation, the strategy space of each user r is $S_r = [0, \infty)$. It may be reasonable to constrain this further for a given user, in which case the strategy space is accordingly modified. This reveals an important distinction between the strategy space and the payoff function: *joint* effects of actions are reflected in the payoff function (e.g., congestion); *individual* constraints on actions are reflected in the strategy space (e.g., an individual resource constraint).

3.1.3 Discussion

Non-atomic games. In the preceding basic model of a static game, the set of players is a finite set. At times it can be useful to relax this assumption and consider games with a continuum of players. Such models are referred to as *non-atomic games* and are a straightforward generalization of the models considered here: nonatomic games can be viewed as limiting models of finite games as the number of players increase. Often in such limits, the impact of any one agent on the other agents becomes negligible, though the impact of a mass of agents can still be significant. For example, this can capture the notion that under some pricing mechanisms, as the number of users increases, each user effectively becomes a price taker.

Simultaneity. One way to think about a static game is that it is modeling a situation in which the players simultaneously choose their

actions, without knowledge of the choice made by any other player. After all players have chosen their actions, they each receive the corresponding payoff. In principle, the choice of actions need not be truly simultaneous; it must simply be the case that each agent chooses their action before finding out about the choice of any other.

Rationality and selfishness. A standard assumption in game theory is that each player behaves *rationally*; that is, each player chooses her actions individually to maximize her own payoff. (See also the discussion in Section 2.1.) To emphasize this point, this type of model is sometimes referred to as a *non-cooperative* game and the agents are referred to as being “selfish.” However, this terminology should not be taken to mean that agents in a game will never cooperate and must be directly competing with each other. It simply means that any cooperation emerges as a result of their individual decisions; whether this happens or not depends on the payoffs assigned to the players. In particular, the literature sometimes confusingly distinguishes between “selfish” players, “altruistic” players (who seek only to help other players), and “malicious” or “adversarial” players (who seek only to hurt other players). To a game theorist, all such players are rational payoff optimizers — the difference is in the form of each player’s payoff.

Dynamics. The model as we have defined it appears to preclude *dynamic* models of games, where agents take decisions over time. Some simple dynamics can be captured even within this model. For example, consider a variation of the peering game in Example 3.1, where the two ISPs make decision about where to route traffic on each day for a week and receive a payoff that is the sum of their daily payoffs. Though they interact over a week, this can still be modeled as a static game, provided that the two ISPs commit to a routing choice for the entire week before seeing the choice of the other player.

More generally, however, it seems reasonable to allow each ISP to observe the other ISP’s routing choice on the previous day. In this case, an ISP should rationally optimize taking this information into account on the current day. Such considerations are not captured in a static game model, and are instead covered by generalizations of this model referred to as *repeated games*. See Section 3.6 for details.

3.2 Dominant Strategies

For a game theoretic model to be useful, one needs to be able to make predictions about the outcome of the game. In game theory such predictions are based on different *solution concepts* that specify properties an outcome of the game must satisfy. One of the simplest solution concepts is that of a dominant strategy, which we discuss here.

3.2.1 Methodology

In a game, player r is said to have a *strict dominant strategy* s_r^* if she prefers that action regardless of the choice of the other players, i.e., there holds:

$$\Pi_r(s_r^*, \mathbf{s}_{-r}) > \Pi_r(s_r, \mathbf{s}_{-r}) \quad \text{for all } s_r \in S_r, \mathbf{s}_{-r} \in S_{-r}. \quad (3.2)$$

If the strict inequality in (3.2) is replaced with a weak inequality, then the corresponding strategy is referred to as a *weak dominant strategy* for player r . If dominant strategies exist for every player in a game then such a choice of strategies by each player is said to be a *dominant strategy equilibrium*.

3.2.2 Example

The following example illustrates a dominant strategy equilibrium for the peering game in Example 3.1.

Example 3.3 (Peering (continued)). Observe that in the peering game *regardless* of the action taken by the other ISP, each ISP is always strictly better off choosing the nearest exit, i.e., the nearest exist is a dominant strategy. Note however, when both ISPs follow that approach, they leave themselves with payoffs uniformly worse than those obtained if both had chosen the further exit. (Classically, this result is a particular instance of the “prisoner’s dilemma.”; see the endnotes for references.)

Although simplified, this model captures a common phenomenon in modern interdomain routing: the *hot potato effect* where all providers try to ensure traffic leaves their networks as quickly as possible. This

simple model shows that this is a natural equilibrium consequence of a cost minimization on the part of the providers, even though it may lead to highly inefficient outcomes.

3.2.3 Discussion

Nonexistence. If a game has a dominant strategy equilibrium, then this is a natural prediction for the outcome of a game. In particular, for such an equilibrium to arise, it only requires that agents are rational (i.e., they seek to optimize their payoffs), and that they know their own strategy set and payoff function. They do not need to know the other agents' actions, or even other agents' payoff functions. Thus, aside from the issues of rationality discussed in the previous section, the dominant strategy equilibrium is a compelling solution concept.

However, there is one significant issue with dominant strategies: they may not exist. The definition of a dominant strategy is quite strong, because it requires a single action to be optimal over *all* possible profiles of others' actions. It is not surprising, therefore, that most games don't possess dominant strategy equilibria. For example, for most reasonable choices of utilities and latency functions, the resource sharing game in Example 3.2 does not have dominant strategies.

Dominated strategies. Even if a game does not have a dominant strategy equilibrium, one can still use the notion of dominance to remove some strategies from consideration. For example if a player r is always strictly better off choosing action s_r over action s'_r , regardless of the choices of the other players, then player r will rationally never choose s'_r . In this case, we say that s'_r is *strictly dominated* by s_r . Taking this a bit further, it would seem reasonable for other players to also assume that r would never choose s'_r , i.e., this action could be eliminated from the game. Once all dominated actions are removed, the new game might have a dominant strategy equilibrium, even though the original did not.

Note, however, that justifying this equilibrium requires more of the agents than just rationality. Specifically, they need knowledge of the payoffs of the other users as well as knowledge of the rationality of the other agents; the reasonableness of such assumptions depend highly

on the setting one is considering. For example, one common motivation for applying game theoretic modeling to networking is that one is interested in a setting in which a centralized solution is not feasible due to the lack of knowledge of every player's payoffs at a centralized controller. Assuming each player to have this knowledge themselves seems difficult to justify.

Iterated elimination of dominated strategies. In applying dominance, one does not need to stop with one round of elimination. For example, suppose that we first eliminate a strictly dominated strategy s'_r of agent r . Suppose also that in the resulting game, a strategy s'_q of another agent q becomes strictly dominated. Following the same logic as before, it would appear reasonable to also eliminate s'_q from consideration. Repeating this process until no dominated strategies remain is known as *iterated elimination of strictly dominated strategies*.

Note, however, for agent r to believe player q would never play s'_q , agent r must first *know* that q believes r would never play s'_r . This requires one more "level" of knowledge on the part of agent r , i.e., r needs to know that q knows of r 's payoffs and rationality. Repeated indefinitely, we obtain the notion of *common knowledge*: informally, a fact is common knowledge if all expressions of the form " r knows that q knows that t knows that ... knows the fact" are true. A crucial element of game theory is that *it is common knowledge that all players are rational*; it is this assumption that is typically used to justify iterated elimination of strictly dominated strategies. (Note that this process also requires the structure of the game to be common knowledge.) We note that this is not a benign assumption, and has some surprising consequences; see the endnotes for details.

If the iterated elimination process concludes with a unique strategy remaining for each player, then this can be used as a solution to the game. Each round of elimination yields a "stronger" solution concept that enlarges the set of games we can meaningfully "solve." However, justifying each round of elimination requires one more level of knowledge on the part of the agents. Even with this assumption and allowing for an arbitrary number of rounds of elimination, many games of interest still do not yield solutions. Nash equilibrium, discussed next, gives us a solution concept that is more broadly applicable.

3.3 Nash Equilibria

By far the most common solution concept used in game theory is the notion of a *Nash equilibrium*. One reason for this is that Nash equilibria exist for many games of interest, though they need not exist for all games or be unique. Loosely, a Nash equilibrium represents a “stable” profile of actions across the users, in the sense that no one user ever can improve her payoff by deviating from a Nash equilibrium action, provided that all other users keep their equilibrium actions fixed.

3.3.1 Methodology

Given a game G , if player r knew that the other players were choosing actions $\mathbf{s}_{-r} \in S_{-r}$, then to maximize her payoff, she should choose an action from the set

$$\mathcal{B}_r(\mathbf{s}_{-r}) = \arg \max_{s_r \in S_r} \Pi_r(s_r, \mathbf{s}_{-r}). \quad (3.3)$$

The set-valued function \mathcal{B}_r is referred to as player r 's *best-response correspondence*²; any element of the set $\mathcal{B}_r(\mathbf{s}_{-r})$ is called a *best response* of player r to the other player's actions. A Nash equilibrium is defined in terms of these correspondences as follows.

Definition 3.2. An outcome \mathbf{s}^e is a *Nash equilibrium* of a game if and only if $s_r^e \in \mathcal{B}_r(s_{-r}^e)$ for all r .

Thus at a Nash equilibrium, each player's action is a best response to the other players' actions. Clearly, if each player r has a (weak or strict) dominant strategy, it will be a best response to any \mathbf{s}_{-r} . Hence, a dominant strategy equilibrium is also a Nash equilibrium. However, in general, for player r to select an action s_r in a Nash equilibrium, it need not be dominant; it only needs to be a best response to the *particular* choice of actions selected by the other players in the equilibrium.

Not all games have a Nash equilibrium, while others may have multiple Nash equilibria. Establishing the existence of a Nash equilibrium and characterizing its uniqueness is then a basic question. (It is

²A correspondence is simply another name for a set-valued function.

a well known fact that existence of Nash equilibria can typically be guaranteed if players are allowed to *randomize* over available actions; however, we defer discussion of this point to Section 3.4 on mixed strategies.)

Broadly, there are two approaches to proving existence of Nash equilibria in general. The first approach relies on the fact that a Nash equilibrium is a *fixed point* of the best response correspondence, and thus certain fixed point theorems can be used to establish existence. The second approach is constructive, and depends on specific structure in the game that can be exploited to verify an equilibrium exists. To illustrate these approaches we discuss *concave games* (an example of the first type) and *potential games* (an example of the second type).

We start with *concave games*, which are defined in terms of convexity properties of the payoffs and strategy spaces.

Definition 3.3. A strategic form game $(R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ is a *concave game* if for each player $r \in R$:

- (1) S_r is a nonempty compact, convex subset of \mathbb{R}^n (for some n);
 - (2) $\Pi_r(s_r, \mathbf{s}_{-r})$ is continuous in \mathbf{s} for all feasible \mathbf{s}_{-r} ; and
 - (3) $\Pi_r(s_r, \mathbf{s}_{-r})$ is a concave function of s_r for all feasible \mathbf{s}_{-r} .
-

Theorem 3.1. Every concave game has at least one Nash equilibrium.

Proof. We only sketch the proof here. It relies on characterizing Nash equilibria in terms of the *joint* best-response correspondence given by $\mathcal{B}(\mathbf{s}) = \prod_r \mathcal{B}_r(\mathbf{s})$, which maps an outcome \mathbf{s} into the Cartesian product of the individual best responses to this outcome. An outcome \mathbf{s}^e is a Nash equilibrium if and only if it is a fixed point of \mathcal{B} , i.e., if $\mathbf{s}^e \in \mathcal{B}(\mathbf{s}^e)$. Kakutani's fixed point theorem gives conditions for a set-valued function to have such a fixed point. For a concave game, it can be shown that \mathcal{B} satisfies the conditions of this theorem. \square

Theorem 3.1 can also be generalized to the class of *quasi-concave games*. These are defined as in Definition 3.3, except the third condition

is replaced by specifying that $\Pi_r(s_r, \mathbf{s}_{-r})$ be a quasi-concave function of s_r for all feasible \mathbf{s}_{-r} .

Potential games are a second class of games where Nash equilibrium existence can be guaranteed. This class is not defined in terms of convexity properties, but instead relies on the payoffs of each player being “aligned” toward a common objective.

Definition 3.4. A strategic game $G = (R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ is a potential game if there exists a function $P : S \mapsto \mathbb{R}$ such that for all $r \in R$, for all $\mathbf{s}_{-r} \in S_{-r}$ and all $s_r, t_r \in S_r$,

$$\Pi_r(s_r, \mathbf{s}_{-r}) - \Pi_r(t_r, \mathbf{s}_{-r}) = P(s_r, \mathbf{s}_{-r}) - P(t_r, \mathbf{s}_{-r}).$$

(This class of games is sometimes called an *exact* potential game to distinguish it from other notions of potentials that have also been studied.)

The function P in this definition is called a *potential function* for G . This definition implies that if G is a potential game, then the change in the payoff each agent r sees from any unilateral deviation from the profile (s_r, \mathbf{s}_{-r}) is exactly the same as the change that agent would see if her payoff was replaced by the common potential function.

Potential functions allow a direct, constructive argument to establish existence of Nash equilibria. In particular, if a strategy profile $\mathbf{s} \in S$ maximizes $P(\mathbf{s})$ over S , then the maximizing profile must be a Nash equilibrium of G : if it were not, an agent could improve her payoff and thus also increase the value of P . This is summarized in the following theorem.

Theorem 3.2. If G is a potential game with potential P and P attains a maximum over S , then G has a Nash equilibrium.

For example, this theorem applies to any game with finite strategy sets, or to any game where P is continuous and S is compact. In general, the converse of the preceding result is not true, i.e., a potential game may have a Nash equilibrium that is not a maximum of P . However, if P is continuously differentiable and S is compact, all Nash equilibria are

stationary points of P (i.e., the first order conditions for maximization of P over S are satisfied). In this case studying the Nash equilibria of G reduces to simply studying the optimization problem $\max_{\mathbf{s} \in S} P(\mathbf{s})$. It follows that if in addition P is strictly concave over S , then G will have a unique Nash equilibrium.

3.3.2 Examples

Next we illustrate existence of Nash equilibria for two variations of the resource allocation game in Example 3.2, corresponding to a concave game and a potential game, respectively.

Example 3.4 (Allocation of a congestible resource as a concave game). Consider the resource allocation game of Example 3.2. Assume that each player r 's utility function U_r is concave, continuously differentiable, and strictly increasing. Further, assume that the function ℓ is convex, continuously differentiable, and strictly increasing. Then it is easily verified that each player's payoff (3.1) is a concave function of the player's own action, and continuous in all the other players' actions. Further, since $U_r(x_r)$ can grow at most linearly in x_r , while $x_r \ell(\sum_s x_s)$ grows at least quadratically in x_r , it follows that we can place an upper bound M on the maximum possible rate that any user would wish to send; this ensures the strategy space of each user is the compact, convex subset $[0, M]$. (We can choose M as a rate such that if $x_r \geq M$, then player r 's payoff is negative regardless of the rates chosen by other users.) Thus, this is a concave game and by Theorem 3.1, it has at least one Nash equilibrium.

Note that since the game is concave, first-order conditions suffice to establish optimality of a best response for a user. In particular, if the other users are sending at rates \mathbf{x}_{-r} , then the first-order conditions for optimality for user r are:

$$U_r'(x_r) - \ell\left(\sum_s x_s\right) - x_r \ell'\left(\sum_s x_s\right) = 0,$$

if $x_r > 0$, and $U_r'(0) \leq \ell(\sum_s x_s)$ if $x_r = 0$. It is straightforward to show that under our assumptions on U_r and ℓ , these equations have a unique

joint solution; we leave the details as an exercise for the reader. Thus in this setting, there exists a *unique* Nash equilibrium.

Example 3.5 (Allocation of a congestible resource as a potential game). As a second example consider a variation in the resource allocation game of Example 3.2 in which the latency function $\ell(y)$ models the total latency cost imposed on a user (instead of a cost per unit resource). In this case each user's payoff becomes

$$\Pi_r(x_r, \mathbf{x}_{-r}) = U_r(x_r) - \ell\left(\sum_s x_s\right). \quad (3.4)$$

It is then straightforward to show that the following function is a potential for the resulting game:

$$P(\mathbf{x}) = \sum_r U_r(x_r) - \ell\left(\sum_s x_s\right). \quad (3.5)$$

It follows from Theorem 3.2 that this game has at least one Nash equilibrium (provided that $P(\mathbf{x})$ is a continuous function).

3.3.3 Discussion

As noted in the previous section, obtaining a solution concept that applies to a larger class of games requires more work to justify. Justifying a dominant strategy equilibrium requires only rationality. Justifying iterated elimination of dominated strategies requires common knowledge that all players are rational (i.e., payoff maximizers).

In general, justifying a Nash equilibrium requires more than this. An agent's action in a Nash equilibria is a rational response to the equilibrium profile of the other users, but why should the other agents' choose this particular profile? Moreover, how does the agent in question know other agents will choose this profile? Justification can be especially troublesome if there are multiple possible Nash equilibria.

There are several possible ways a Nash equilibrium is justified, depending in part on how game theory is being applied. A useful

dichotomy is given by the “semantic” and “economic” approaches discussed in the Introduction. When a semantic approach is taken, the reason for applying a given solution concept is often simply that the concept reflects the behavior of the underlying protocol or algorithm. On the other hand, when an economic approach to game theory is taken, the motivation for a given solution concept becomes more challenging, since the solution concept is being used to predict the behavior of rational players with their own incentives.

In this section we provide some discussion of possible justifications for Nash equilibrium.

Nash equilibrium as a self-enforcing outcome. Our first interpretation is based on considering a “one-shot” game in which the players meet and play the game in question exactly once. In this context, what might lead players to play a Nash equilibrium? At the very least, Nash equilibrium also requires that rationality of the players and the structure of the game are common knowledge. However, as noted above, Nash equilibrium requires some additional justification. One common approach is to assume a non-binding agreement on the player’s actions before playing the game. In a communication networking setting, such an agreement could be in the form of a non-binding standard that the agents agree to follow. If the players agreed on a Nash equilibrium, then this agreement would be *self-enforcing* and so no player would have an incentive to break it (e.g., use devices or protocols that are not compliant with the agreed standard).

Nash equilibrium as the outcome of long-run learning. An alternative interpretation of a Nash equilibrium is to assume that the game is actually modeling a situation in which each player plays the same game many times, and “learns” how to play over the long run. Over time, each player is assumed to be able to experiment with different actions to try to improve her payoff in the game. If such a process reaches a “steady-state” (where no player wants to change, given what others are playing), this will necessarily be a Nash equilibrium. With this interpretation, a player need not have full information about the other players’ payoffs or make assumptions about their rationality, but may learn enough by repeatedly playing the game to reach a Nash equilibrium.

We have been somewhat vague about exactly what learning process a player uses. Indeed, the interaction between equilibrium concepts and the dynamics which lead to them has become a fruitful area of research in game theory in the recent past. As one example, we discuss a particular learning dynamic referred to as *simultaneous best response dynamics* (also known as *best reply dynamics* or *myopic best response dynamics*).

Consider a strategic game $G = (R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ which the R players repeatedly play in a sequence of rounds. We assume that initially, each player r only knows her own payoff Π_r and her own strategy set S_r . Let $s_r^{(n)}$ denote the action of player r in the n th round. After a given round n each player r observes the actions of the other players $\mathbf{s}_{-r}^{(n)}$ and her own payoff $\Pi_r(s_r^{(n)}, \mathbf{s}_{-r}^{(n)})$. Under a simultaneous best response dynamic, each player r then updates her strategy for round $n + 1$ as follows:

$$s_r^{(n+1)} = \mathcal{B}_r(\mathbf{s}_{-r}^{(n)}),$$

i.e., each player's action in round $n + 1$ is a best response to her opponents action in round n . (If an agent has multiple best responses, then it is also useful to add the restriction that it will not change its response unless it improves its payoff.) Here, we assumed that all player's choose their best reply in each round. A related dynamic is to assume *sequential best responses* under which in a given round exactly one player updates; the ordering of the players can be arbitrary with the restriction that every player updates within some finite time.

Clearly if either of these best response dynamics converge, then the limit point must be a Nash equilibrium. Unfortunately, these dynamics are not guaranteed to converge in all games for which a Nash equilibrium exists. Showing convergence can either be done directly (e.g., by using contraction mapping arguments) or by appealing to various results in the literature which identify classes of games for which these dynamics are known to converge. For example, sequential best response dynamics must converge in a potential game with finite strategy sets, since at each round the potential function must be increasing and it only takes on a finite number of values.

In addition to best response dynamics, there are many other learning dynamics that have been studied for games. However, a potential weakness with such an approach is that it does not provide a convincing story as to *why* agents would all choose a given learning rule. For example, if two agents are playing a game repeatedly, then an agent may have an incentive to deviate from the learning rule in a given round to “fool” her opponent and receive a larger payoff later. Indeed, if we account for the fact that players should both track the history of the game and forecast how they expect their opponents to behave, then we are in the realm of dynamic games; or more precisely, a subclass of dynamic games known as *repeated games* (see Section 3.6).

One way to preclude such strategizing from round to round is assume that in each round a player faces a different set of opponents drawn from a large population. In such a setting, best response updates do not make sense, but other learning dynamics can be applied and shown to converge to a Nash equilibrium (or closely related concepts) for certain classes of games. Of course assuming a large population of potential opponents does not make sense for every application, but in cases where users interact as part of a large network (e.g., the Internet) such an approach may have merit.

Nash equilibrium as an engineering goal. Dynamic adjustment processes are often at the heart of the semantic approach to applying game theory to networks. In particular, in protocol design settings where a Nash equilibrium is actually the desired “stable point” of the system, one can interpret the learning rule as part of the protocol description. When applied to the design of new protocols, the question then becomes how to design an agent’s payoffs and the corresponding learning rule so that it converges to a desirable equilibrium. In this context, many of the objections regarding justification of Nash equilibrium go away; in particular, in this setting it is typically assumed there is no risk that agents (e.g., devices, software processes) will “deviate” from the prescribed learning protocol.

3.4 Mixed Strategies

So far we have assumed that each agent in a game deterministically chooses exactly one action from her strategy set. Such a choice

is referred to as a *pure strategy*. At times, it is advantageous to allow agents to instead randomize their decision, resulting in what is called a *mixed strategy*. In particular, allowing mixed strategies ensures that (under weak conditions) a game possesses at least one Nash equilibrium.

3.4.1 Methodology

Mixed strategies are formally defined as follows.

Definition 3.5. A *mixed strategy* for player r with strategy set S_r is a probability distribution σ_r on S_r .

For the case where S_r is finite, then for any action $s_r \in S_r$, $\sigma_r(s_r)$ can be interpreted as the probability that player r plays action s_r .

Let $\Delta(S_r)$ be the set of all probability distributions over S_r . If agents in a game $G = (R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ can play mixed strategies, then they can be viewed as choosing pure strategies in the game $G' = (R, \{\Delta(S_r)\}_{r \in R}, \{\tilde{\Pi}_r\}_{r \in R})$, which is referred to as the *mixed extension* of G . Here the payoffs $\tilde{\Pi}_r$ in the mixed extension are interpreted as *expectations* over the joint distribution on actions induced by the mixed strategy profile. A *mixed Nash equilibrium* of a game is then defined to be a Nash equilibrium of the mixed extension.

Consider the special case where S_r is a finite set for each player r . Note that in this case, the mixed extension is a game where the strategy spaces are compact and convex. The payoff of each player r is given by

$$\tilde{\Pi}_r(\sigma_r, \sigma_{-r}) = \sum_{\mathbf{s} \in S} \Pi_r(\mathbf{s}) \prod_{q \in R} \sigma_q(s_q), \quad (3.6)$$

which is continuous in the entire strategy vector, and concave (in fact, linear) in a player's own strategy. As a result, the mixed extension is a concave game; and so a Nash equilibrium exists, by Theorem 3.1. Thus we have the following theorem.

Theorem 3.3. Suppose that $G = (R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ is a game where each S_r is a finite set. Then there exists a mixed Nash equilibrium of this game.

This result is one of the most celebrated insights in game theory, due to John Nash (hence the term Nash equilibrium).

When each player has a finite strategy set, the payoff in (3.6) satisfies

$$\tilde{\Pi}_r(\boldsymbol{\sigma}_r, \boldsymbol{\sigma}_{-r}) = \sum_{s_r \in S_r} \sigma_r(s_r) \tilde{\Pi}_r(\delta(s_r), \boldsymbol{\sigma}_{-r}),$$

where $\delta(s_r)$ indicates the (degenerate) mixed strategy which selects s_r with probability one (i.e., the pure strategy s_r). It follows that in a mixed Nash equilibrium, player r must assign positive probability only to strategies that are a best response to $\boldsymbol{\sigma}_{-r}$, and thus player r will be indifferent between each of these strategies. It can be shown that this is also a sufficient condition for a strategy profile to be a Nash equilibrium. This provides a useful approach for calculating mixed equilibria, which we summarize in the following lemma.

Lemma 3.4. Suppose that $G = (R, \{S_r\}_{r \in R}, \{\Pi_r\}_{r \in R})$ is a game where each S_r is a finite set. Then $\boldsymbol{\sigma}$ is a mixed Nash equilibrium if and only if for every $r \in R$, σ_r assigns positive probability only to pure strategies that are best responses to $\boldsymbol{\sigma}_{-r}$.

3.4.2 Example

The next example illustrates how randomized strategies might arise in the context of accessing a shared medium.

Example 3.6 (Random access game). Suppose that R nodes are accessing a shared medium, such as a wireless channel. Consider a game in which each node has two pure strategies: transmit (denoted T) and not transmit (denoted NT). Each transmission incurs a “cost” of $-c$. If a node transmits and no other node does so, the transmission is successful, and the node receives a payoff of $1 - c$; otherwise, if at least one other node transmits at the same time, a collision occurs and the payoff is $-c$. Finally, if a node does not transmit, the payoff is zero.

It can be seen that this game has R pure strategy Nash equilibria, each corresponding to one node transmitting and all other nodes not

transmitting. Such equilibria clearly favor the user who is transmitting, and furthermore, do not provide an adequate explanation of how the agents select among these equilibria. By turning to mixed strategies, we can uncover an equilibrium that is symmetric among the players.

In particular, this game has a symmetric mixed strategy equilibrium in which each agent transmits independently with probability p (i.e., $\sigma_r(T) = p$ and $\sigma_r(NT) = 1 - p$ for all r); in other words, each node uses a random access protocol. From Lemma 3.4, it follows that a necessary and sufficient condition for such an equilibrium to exist is that

$$(1 - c)(1 - p)^{R-1} - c(1 - (1 - p)^{R-1}) = 0.$$

Solving for p yields $p = 1 - c^{\frac{1}{R-1}}$. (Note that if $c = (1 - 1/R)^{R-1}$, then $p = 1/R$, which maximizes the total throughput of the players among all symmetric strategies.)

3.4.3 Discussion

Interpreting randomization. A mixed Nash equilibrium essentially implies that players use randomization as a fundamental part of their strategy. As with the definition of Nash equilibrium itself, a variety of justifications are possible for the use of randomization. In particular, in the context of an engineering system, it is often natural to introduce randomization. Indeed, randomization is part of many common networking protocols, e.g., in a random access protocol as in the previous example.

However, when modeling economic agents (e.g., people, firms), interpreting such randomization is a subject of debate. For example, behavioral studies have shown that people rarely randomize. Moreover if agents meet and play a game once, any randomization on their parts is never observed. Finally, in a mixed strategy equilibrium, players are indifferent over the actions they assign positive probabilities to. Why should they assign these probabilities in a way that satisfies the equilibrium conditions?

Instead of assuming that agents internally randomize when deciding on their action, several alternative justifications for this concept have

also been developed. One of these, known as a *purification* approach, is to view any randomization as a modeling tool that arises due to small imperfections in our model of the “true game.” (See the endnotes for details.) An alternative justification is to consider a setting with a large number of potential agents who may play the game. Each agent has a pure strategy; the randomization reflects the distribution of strategies across the population. (Of course, the latter explanation still begs the question of how this distribution arose in the first place.) As with other elements of game theory, justifying mixed strategies is more challenging in economic applications than in semantic applications of the theory.

3.5 Bayesian Games

In game theory, *complete information* refers to any setting where the structure and rules of the game, the payoff functions of the players, and the rationality of all players, are all common knowledge. All the models we have studied thus far fall in this class. This is a strong assumption in many cases and it may often be more reasonable to assume that agents only have partial information about the game. For example, when multiple agents are competing for a common resource, an agent may know its own value of the resource, but not have exact knowledge of the other agents’ valuations.

Bayesian games or *games with incomplete information* provide a framework for modeling such situations. *Incomplete information* refers to a setting where some element of the structure of the game — typically the payoffs of other players — are *not* common knowledge. A significant challenge then arises: how should agents reason about each other? The economist John Harsanyi had a key insight in addressing this issue: uncertainty can be modeled by introducing a set of possible “types” for each player, which in turn influence the players’ payoffs. Players know their own type, but only have a probabilistic belief about the types of the other players. In this way, players can reason about the game through a structured model of their uncertainty about other players.

3.5.1 Methodology

Formally, a Bayesian game is defined as follows.

Definition 3.6. A static Bayesian game G is a 5-tuple $(R, \{S_r\}_{r \in R}, \{\Theta_r\}_{r \in R}, F, \{\Pi_r\}_{r \in R})$ where:

- (1) R is a finite set of players;
 - (2) for all $r \in R$, S_r is the *strategy set* of player r ;
 - (3) for all $r \in R$, Θ_r is the *type set* of player r ;
 - (4) $F : \prod_{r=1}^R \Theta_r \mapsto [0, 1]$ is a joint probability distribution over the type space; and
 - (5) for all $r \in R$, $\Pi_r : \prod_{s=1}^R S_s \times \prod_{s=1}^R \Theta_s \mapsto \mathbb{R}$ is the *payoff function* of user r .
-

The set of players, their strategy sets, and the payoffs have the same interpretation here as in a static game with complete information. The new ingredients here are the type set for each player r and the probability distribution F over the type space $\Theta = \prod_{r=1}^R \Theta_r$. This probability distribution is sometimes referred to as a *common prior*.

An alternative way of defining Bayesian games is by introducing a set of “states of the world,” Ω . One such state is chosen randomly by “nature” and the players’ types are given by functions that map the state of the world to their type. A player’s payoff then depends on the chosen state and the action profile. In this setting, the distribution of nature’s decision is the common prior.

In a Bayesian game, all players observe their *own* type, but do not observe *others’* types. However, using the prior they can formulate a conditional distribution over other players’ types, given their own privately observed type. One can view the game as if “nature” first draws a profile of types according to the common prior and reveals to each player her own type $\theta_r \in \Theta_r$, but not the type chosen for any other player. Each player’s payoff is also now parameterized by the chosen type profile as well as by the action chosen for each player r .

How should players reason about such a game? For a player to determine her own payoff, she needs to know how other players will behave;

and to determine how other players will behave, she needs a model of their strategic choices for *every* type realization that has nonzero probability of occurring. As a consequence, a strategy for player r in a Bayesian game is actually a *function* $a_r : \Theta_r \mapsto S_r$, which specifies the action chosen for each possible type. Mixed strategies can again be defined as a probability distribution over such pure strategies. To simplify our presentation, we focus only on pure strategies in the following.

Next we turn to solution concepts. A dominant strategy equilibrium for a Bayesian game can be defined in an analogous manner as before; only here a player is required to have a dominant action for every possible profile of types. Also, as before, this is often too strong of a solution concept to be useful for many games.

Instead, the most common solution concept used for Bayesian games is *Bayesian–Nash equilibrium*, which naturally generalizes Nash equilibrium to the Bayesian setting. Once again we define equilibria in terms of best response correspondences. However, now in our definition we need to account for the incomplete information in the game. Specifically, conditioned on a player knowing her own type, each player in a Bayesian game is assumed to maximize their expected utility, where the expectation is taken over the remaining uncertainty in the types of the other players. Given a profile of strategies \mathbf{a}_{-r} for the other players, and a type θ_r for player r , this leads us to define player r 's best response correspondence conditioned on her type, as follows:

$$\mathcal{B}_r(\mathbf{a}_{-r}|\theta_r) = \arg \max_{s_r \in S_r} E_{\theta_{-r}}[\Pi_r(s_r, \mathbf{a}_{-r}(\theta_{-r}), \boldsymbol{\theta})|\theta_r].$$

Here, the expectation is taken over θ_{-r} , the types of the remaining players, conditioned on θ_r .

We can likewise represent the (unconditioned) best response strategies a_r of player r as

$$\mathcal{B}_r(\mathbf{a}_{-r}) = \arg \max_{a_r: \Theta_r \mapsto S_r} E_{\boldsymbol{\theta}}[\Pi_r(a_r(\theta_r), \mathbf{a}_{-r}(\theta_{-r}), \boldsymbol{\theta})],$$

where the quantity on the right is the agent's *ex ante* expected utility, i.e., the utility that agent expects to get before seeing her type. Note that

$$\mathcal{B}_r(\mathbf{a}_{-r}) = E_{\theta_r}[\mathcal{B}_r(\mathbf{a}_{-r}|\theta_r)].$$

Definition 3.7. A strategy profile \mathbf{a}^e is a *Bayesian–Nash equilibrium* of a game if and only if $a_r^e \in \mathcal{B}_r(\mathbf{a}_{-r}^e)$ for all r .

Such equilibria can be viewed as Nash equilibria of an “expanded game” in which each player’s strategy set is the set of all functions $a_r : \Theta_r \mapsto S_r$ and their payoffs are given by their *ex ante* expected utility.

As in the case of games with complete information, various results are known that help to establish the existence of Bayesian–Nash equilibrium. For example, the following result is analogous to Theorem 3.1.

Theorem 3.5. Consider a Bayesian game in which the strategy spaces and type sets are compact subsets of \mathbb{R} and the payoff functions are continuous and concave in a player’s own strategy. Then a pure strategy Bayesian–Nash equilibrium exists.

Likewise, Nash’s theorem for mixed strategies naturally generalizes as in the following.

Theorem 3.6. Mixed Bayesian–Nash equilibria exist for all finite games.

3.5.2 Example

We give two examples of Bayesian games.

Example 3.7 (First price auction). An object is to be assigned to one of two agents via a *first price auction*. In this mechanism, each player submits a non-negative bid b_r ; the object is awarded to the agent with the highest bid. The winning bidder must pay her bid, while the losing bidder pays nothing. (In the case of a tie, assume no one is awarded the object, nor pays anything.) Each agent r has a private value for the object of θ_r . We assume an independent uniform prior on the valuations; that is, we assume that the values are generated independently chosen according to the uniform distribution on $[0, 1]$. This is an example of an *independent private values* model.

This situation can be readily modeled as a Bayesian game, where the strategy set for each player r corresponds to the choice of possible bids and the type of each player corresponds to their private value for the object. Thus, the type set Θ_r of each player is simply the interval $[0, 1]$, and the common prior is simply the product of two uniform measures on $[0, 1] \times [0, 1]$. The payoff function of a user in this game is then given by

$$\Pi_r(b_r, b_{-r}, \theta_r) = \begin{cases} \theta_r - b_r, & \text{if } b_r > b_{-r}, \\ 0, & \text{otherwise,} \end{cases} \quad (3.7)$$

which in this case does not depend on the type of the other agent.

We now show that a Bayesian–Nash equilibrium for this game is for each player to bid half their value, i.e., $b_r = \theta_r/2$. To see that this is an equilibrium, consider agent 1’s best response given that agent 2 is bidding half her value. Player 1 will then win the auction if her bid b_1 is larger than $\theta_2/2$. From the assumed uniform prior distribution, the probability that this occurs is given by $\min(2b_1, 1)$, in which case player 1 receives a payoff of $\theta_1 - b_1$. It follows that agent 1’s best response given type θ_1 is the value of $b_1 \geq 0$ which maximizes

$$E_{\theta_2}[\Pi_1(b_1, \theta_2/2, \theta_1)|\theta_1] = (\theta_1 - b_1) \min(2b_1, 1).$$

The unique maximizer of the right-hand side is $b_1 = \theta_1/2$. In other words, agent 1’s best response is to also bid half her value. Hence, from symmetry it follows that this must be a Bayesian–Nash equilibrium.

Note that bidders in a first price auction are incentivized to *shade* their bids downward, so that they make a positive profit when they win the auction. (For more on auctions, see Section 5.4.)

Example 3.8 (The market for lemons revisited). In the previous example, a player’s payoff depended only on her own type and so each player was uncertain about the payoff of the other player in the game (hence the term “private values”). Bayesian games can also be used to model other forms on uncertainty. We illustrate this here by developing a Bayesian game to model the “market for lemons” example discussed in Section 2.6.

Recall that in this example a buyer wants to purchase a single unit of a good that the seller values at v and the buyer at αv , for $\alpha > 1$. The seller's value v is a uniform random variable on $[0, 1]$ and is observed by the seller and not the buyer. A price p for the good is set exogenously in a market, and the question considered was at what price p would trade take place. To view this as a Bayesian game, we let the set of agents be the buyer b and the seller s . The seller has two actions: "sell" or "not sell"; the buyer also has two actions: "buy" or "not buy." We can view the seller as having a type v that corresponds to her value, while the buyer has a single type with probability one. Let the seller's payoff be given by

$$\Pi_s(s_s, s_b, v) = \begin{cases} p, & \text{if } s_s = \text{sell and } s_b = \text{buy,} \\ v, & \text{otherwise.} \end{cases}$$

The buyer's payoff is

$$\Pi_b(s_s, s_b, v) = \begin{cases} \alpha v - p, & \text{if } s_s = \text{sell and } s_b = \text{buy,} \\ 0, & \text{otherwise.} \end{cases}$$

This completes our description of the game. Note that in this case, the buyer's payoff depends on the type of the seller and so the buyer *does not know* her own payoff; while the seller knows both her own payoff and that of the buyer. (Settings where all players' payoffs depend on a common source of underlying uncertainty are known as *common value* settings; see Section 5.4 for further discussion.) Following the discussion in Section 2.6, it can be verified that for $\alpha < 2$, a Bayesian-Nash equilibrium of the game is for the buyer to always choose "not buy" so that trade never occurs.

3.5.3 Discussion

The common prior. Perhaps the central objection to Bayesian games is that they rely on the notion of a common prior between agents. Relative to Nash equilibrium, the advantage is that for rational players to select a Bayesian-Nash equilibrium, they no longer need knowledge of other players' payoff functions; instead, they need to know set of possible

types and the common prior in order to determine their expected utility. Moreover, this needs to be common knowledge. The need for a common prior places some significant restrictions on the model that have been disputed. For example, it is not possible for agents in such a game to have “inconsistent” beliefs that are common knowledge.

The model also raises questions about *how* the players arrived at a prior that is common knowledge in the first place. One reaction is to actually use the Bayesian game framework to also model the uncertainty that a player may have about the information available to other players.

For example, in Example 3.8, suppose that the buyer had two types, “informed” and “not informed,” in each case the buyer has the same payoff except in the “informed” case she knows the realization of the seller’s type, while in the “uninformed” case she does not. In this case, the seller would not know if the buyer was informed or not about the value of the good, but rather have a probabilistic belief that the buyer was informed, given by the conditional probability that the buyer had the informed type.

One can continue this conceptual thread and define “higher order” levels of uncertainty, i.e., uncertainty about beliefs about beliefs about beliefs, etc. In principle this can be incorporated into a Bayesian framework, following the same logic as in the previous paragraph. Not surprisingly, however, such models can quickly become intractable; as in much of mathematical modeling, there is a delicate balance between realism and tractability in the model.

Other forms of uncertainty. The previous examples illustrate that conceptually, the Bayesian game framework admits significant flexibility (if at the expense of tractability). As another example, at first glance it may appear that this framework only allows one to capture payoff uncertainty, and not other forms of uncertainty such as uncertainty about the number of agents participating in a game or uncertainty about the actions available to other players. However, by appropriately defining the game, these can also be modeled using a Bayesian game. For example, to model uncertainty about the number of players, one can enlarge a player’s strategy set to include the action “do not participate” and make this action strictly dominant for certain realizations of

a player's type. One can also model variability about the actions that are available by assigning some actions a very low payoff for certain types, so that we can be sure that they would not be used. As before, however, such model complexity is typically greeted with an increase in analytical complexity as well.

3.6 Repeated Games

Thus far, we have focused on “one-shot” models in which agents interact once with each other. More generally, game theory offers a rich modeling framework for *dynamic* interactions as well. In this section, we consider one benchmark model of dynamic games, referred to as *repeated games*. As we shall see, repeated games are conceptually appealing, but also suffer from some analytical pitfalls that limit their practical applicability.

3.6.1 Methodology

Recall the peering game from Example 3.1 in which two ISPs selected to route traffic through either a “near” or “far” peering point. Both selecting “far” led to the socially optimal solution; however, the unique equilibrium of this game was for both to select “near.” This can be attributed to each ISP not accounting for the negative effect, or externality, that their routing choice imposes on the other ISP. (For more on externalities see Chapter 4.)

Suppose now that instead of making a single routing decision, the two ISPs revised their routing decision each day. In such a setting a rational ISP should not just account for the effect of their decision on today's payoff, but also how this will affect future payoffs. Formally, such a setting can be modeled as a *repeated game*, i.e., a model in which the peering game is played repeatedly by the same two opponents in periods $n = 1, 2, \dots$.

At the start of each period the players choose their actions in the current period, given the outcome of the games in all the previous periods. In other words, a player's strategy in such a repeated interaction is a *policy*, which specifies her action at each time n as a function of the history of play up to time n . When choosing such a policy, suppose

that each player's payoff is given by a discounted sum of payoffs in each period:

$$\sum_{n=0}^{\infty} \delta^n \Pi_r(s_r^n, s_{-r}^n),$$

where $\delta \in (0,1)$ is the discount factor, Π_r is the payoff of agent r in a single period and s_r^n is player r 's action in period n . Informally, an "equilibrium" of this game consists of a policy for each player that optimizes their payoff after any history of the game, given the policy chosen by the other player. Such an equilibrium is referred to as a *subgame perfect Nash equilibrium*; see the endnotes for references to a more formal development.

Here we develop some of the main insights of repeated games by way of example. In particular, for the peering example, suppose that ISP 1 adopts the following strategy: it will route traffic to the far peering point in each period as long as ISP 2 does the same; otherwise, it will route all traffic to the near peering point. Given this strategy, let us consider ISP 2's best response. If it routes traffic to the near peering point during period τ , then given the payoff matrix in Example 3.1, it should do this for all periods $n > \tau$. Hence, to determine ISP 2's best response, we only need to determine the first time τ that it routes traffic to the near peering point. For a given choice of τ , its payoff is given by

$$\begin{aligned} & \sum_{n=0}^{\tau-1} \delta^n \cdot (-2) + \delta^\tau \cdot (-1) + \sum_{n=\tau+1}^{\infty} \delta^n \cdot (-4) \\ &= \frac{1}{1-\delta} (-2 + \delta^\tau - 3\delta^{\tau+1}). \end{aligned}$$

For $\delta > 1/3$, this is an increasing function of τ , and thus ISP 2's best response is to always send traffic to the near peering point. It follows that if ISP 2 adopts the same strategy, then this will be a (subgame perfect Nash) equilibrium in which both agents always make the efficient routing choice. In other words, the inefficiency of the prisoner's dilemma is eliminated in this equilibrium.

The strategies in the previous example can be viewed as a type of "punishment" strategy, i.e., in these strategies each agent threatens to

punish the other agent for choosing the “near” route. Such punishment strategies are sometimes also referred to as *trigger strategies*. When the agents do not discount the future “too much,” this threat of punishment gives each agent an incentive to cooperate with their partner.

3.6.2 Discussion

Folk theorem. The preceding discussion suggests, at first glance, that repeated games can perhaps lead to efficiency in settings where the static game might have inefficient outcomes. However, we must be cautious in reaching such a conclusion; in fact, repeated games have many equilibria, some of which are at least as inefficient as any equilibrium of the corresponding static game. For example, in any repeated game, repeating a pure strategy Nash equilibrium of the static game indefinitely is *also* a subgame perfect Nash equilibrium. In the case of the peering game, such an outcome would yield payoffs as low as in the static game.

This example is indicative of a broader result, the so-called “folk theorem” of repeated games. The folk theorem states that for any repeated game, *any* set of feasible *individually rational* payoffs can be supported as an equilibrium, provided that the discounted factor is close enough to 1. “Individually rational” means that the equilibrium payoff of a user must be at least the lowest payoff it can obtain by unilaterally best responding to her opponent’s action, known as the *minimax* payoff:

$$\min_{\mathbf{s}_{-r}} \max_{s_r} \Pi_r(s_r, \mathbf{s}_{-r}).$$

Here player r maximizes her payoff given her opponent’s choice, and in the worst case her opponent is purely adversarial and minimizes player r ’s payoff.

The folk theorem has a constructive proof, and actually yields a pair of strategies that attain any specified payoff vector meeting the constraints. The folk theorem is an example of a result in economics where “anything goes”: essentially, equilibria of repeated games are so numerous as to lack any predictive power. Predictive power can be regained by refining equilibria, i.e., adding progressively stronger

constraints on either the behavior of players or the structure of the equilibrium. However, such constraints typically bring an attendant increase in intractability. As a consequence, while repeated games are a useful modeling tool, they are typically less valuable in providing engineering insight or design guidance for real systems.

Monitoring. In the repeated interaction model we were assuming *perfect public monitoring*, i.e., after each stage both players always observed the action of the other player. Unfortunately, this may not be the case for many models with repeated interactions. In such cases, new subtleties arise. For example, if players observe a noisy version of the other player's actions, then there is a possibility of "false" punishment. Such games with imperfect monitoring have been well studied; by employing more sophisticated strategies, some form of cooperative behavior may still be supported in equilibrium (though the general insight of the folk theorem persists).

Finite repeated games. We have discussed the case where players interact with each other over an infinite time horizon. In practice, most repeated interactions only last a finite number of periods. If the horizon length is fixed and given, the resulting game is a *finitely repeated game*. The equilibria of finitely repeated games can differ significantly from those of infinitely repeated games, depending on the nature of the static game that is repeated. This suggests one should be careful when modeling a particular interaction as an infinitely repeated game, and be sure that it is reasonable that the end of any interaction is far enough in the future that players effectively treat the interaction as indefinite.

3.7 Efficiency Loss

We started this chapter by recalling from Chapter 2 that, for some cases, prices can be used to reach an efficient allocation of resources — provided that users do not anticipate their effect on the price. As we noted, when users do anticipate this effect, the outcome may be different and is better modeled from a game theoretic perspective.

We now turn to a normative question: Given such a game, what loss in efficiency (if any) is incurred by the behavior of fully rational users

who anticipate their effect on the outcome? Answering such a question requires two key choices: first, determining a benchmark metric for measurement of efficiency; and second, choosing an equilibrium concept with which to model the outcome of the game. In this section, we discuss various ways to compare equilibria to the efficient outcome.

3.7.1 Methodology

For simplicity, we focus on resource allocation in a quasilinear environment (as modeled in Section 2.4); we suppress the technical details here.

In particular, consider a game G for allocating one or more resources among agents in a quasilinear environment. The efficient allocation in such a setting is the one which maximizes the total utility across all agents. It appears natural to measure the loss in efficiency for a given outcome of G by the difference between the total utility at the outcome and the total utility at an efficient allocation. Since we are assuming a quasilinear environment, this difference is measured in units of the common currency. It can be convenient to normalize this difference by total utility of an efficient allocation, in which case the efficiency loss is a percentage that does not depend on the units used to measure the common currency. (Note that such a normalization requires the maximum total utility to be positive.) We refer to this (unitless) ratio as the *efficiency loss* at a given outcome of G .

We have discussed various equilibrium concepts for a game; the previous definition can apply to an outcome associated with any such solution concept. Recall that there may be more than one equilibrium, and different equilibria may achieve different utilities. Without a compelling reason to choose one equilibrium over another, one approach for characterizing the performance of a given game is to determine the *worst-case equilibrium*, i.e., the equilibrium with the lowest utility (or largest efficiency loss).

This measure is closely related to the *price of anarchy*, which is defined as the ratio of the total utility achieved by the efficient allocation to the worst (i.e., lowest) total utility achieved by an equilibrium. (Here, we require the total utility of the efficient allocation and the

total utility of any equilibrium to both be positive.) An efficiency loss of 0% corresponds to a price of anarchy of 1; as the efficiency loss increases toward 100%, the price of anarchy increases without bound. Often these concepts are applied to a class of games parameterized by the number of players and their utility functions; in this case, the price of anarchy refers to the the largest price for any game in this class.

3.7.2 Example

We next give an example to illustrate these different measures of efficiency loss.

Example 3.9 (Tragedy of the commons). We consider a special case of Example 3.2 in which R users are sharing a link with capacity C . Let x_r denote the share user r receives. We assume that each user r receives a utility equal to their share, i.e., $U_r(x_r) = x_r$. Furthermore, we assume that the latency function is given by $\ell(y) = y/C$, so that a user's net utility is given by $\Pi_r(x_r, \mathbf{x}_{-r}) = x_r(C - \sum_s x_s)/C$. Our choice of utilities and latency functions is primarily to simplify the following arguments; similar conclusions hold for other choices of these parameters. For the game in Example 3.2, each user chooses x_r to maximize $\Pi_r(x_r, \mathbf{x}_{-r})$. As shown in Example 3.4 this game has a unique Nash equilibrium, which in this case is given by $x_r^e = C/(1 + R)$ for each user r . The social welfare achieved by this equilibrium is given by

$$\sum_r \Pi_r(x_r^e, \mathbf{x}_{-r}^e) = \frac{RC}{(1 + R)^2}.$$

To determine the efficiency loss, suppose instead that a centralized planner knew the utility of each player; then it could instead determine an efficient allocation x'_r . One such allocation is given by $x'_r = C/(2R)$, which results in a social welfare of $C/4$ independent of the number of users. Hence, the (normalized) efficiency loss of the previous Nash equilibrium is $1 - 4R/(1 + R)^2$ and ratio of the efficient total utility to the equilibrium total utility is $(1 + R)^2/(4R)$. If we consider the class of all such games parameterized by R , then the worst-case efficiency loss is 100% and the price of anarchy is unbounded.

This illustrates that when users are allowed to choose their resource allocation in their own self interest the overall social welfare can be greatly reduced, resulting in what is sometimes called a *tragedy of the commons*.

3.7.3 Discussion

Worst-case, average-case, and best-case. Since such measures are worst-case, they may not give an accurate indication of the “average case” that one would find in practice. Note that the price of anarchy is worst-case in two dimensions: both over the family of games under consideration, and over the set of possible equilibria in each game. A related, less pessimistic notion that has also been studied is the “price of stability,” which instead considers the best possible equilibrium in each game.

The efficient benchmark. Because we assumed a quasilinear environment, we were able to rely on the total utility as a benchmark metric for efficiency. In general, however, there may be other metrics that are relevant; e.g., some of the metrics obtained in the fairness discussion in Section 2.3. This is particularly the case when utilities are not being measured in common units, so that total utility is not necessarily the right objective.

Normalization. The efficiency measures we consider here are based on ratio of two utilities. This is only meaningful in cases where the utilities of agents have a common currency of measurement, e.g., delay, throughput, or monetary valuation. As we discussed in Section 2.1, in many economic models this is not the case; a utility is simply a representation of an agent’s underlying preferences over different outcomes. Such a utility function is in general not unique; for example, adding any constant to a utility gives another equivalent utility function. The issue in such cases is that adding a constant to a user’s utility changes the ratio of two utilities used in these efficiency measures, making it unclear what meaning to assign to this ratio. For this reason one must be careful in choosing the correct normalization for efficiency loss. For example, it may make more sense to measure both the utility at the efficient outcome and the utility at the equilibrium outcome as differences from a “reference” outcome.

3.8 Endnotes

Static games. Many of the concepts discussed in this chapter are covered in standard references on game theory such as Fudenberg and Tirole [11], Osborne and Rubinstein [32], or Myerson [28]. Any of these references provide a fuller treatment of advanced topics in game theory, including static and dynamic games with complete and incomplete information. Mas-Colell et al. [24] also contains a treatment of much of this material from a microeconomic perspective.

Dominant strategies. Example 3.1 is equivalent to the well-known “prisoner’s dilemma” game, which is discussed in most introductory books on game theory (e.g., the references described above).

Myerson [28] describes an elegant parable illustrating the pitfalls of a lack of common knowledge reasoning. See also the classic papers of Aumann [2] and Milgrom and Stokey [25].

Nash equilibrium. The notion of a Nash equilibrium for a general static games was introduced by Nash [30], which also contained Theorem 3.3 showing the existence of Nash equilibria in finite games with mixed strategies.

Prior to Nash’s work, the focus of game theory was mainly on two-person zero-sum games: games where one player’s gain is the opponent’s loss. A classic reference of this work is given by Von Neumann and Morgenstern [41].

See Monderer and Shapley [27] for a more extensive discussion of potential games.

The book by Fudenberg and Levine [10] contains a much broader treatment of various learning dynamics for games.

Mixed strategies. The purification approach for justifying mixed strategy equilibrium was given in Harsanyi [15].

Bayesian games. Bayesian games were developed in Harsanyi [14] as was the notion of a Bayesian–Nash equilibrium.

The analysis of the first price auction in Example 3.7 dates back back to the seminal work of Vickrey [40], where it was also shown that the given equilibrium is unique.

Repeated games. Repeated games and the folk theorem are covered in most introductory texts on game theory including those listed previously.

A more extensive treatment of these topics under both perfect and imperfect monitoring is provided in the book by Mailath and Samuelson [23].

The folk theorem for repeated games is so named because this result was part of the “folk knowledge” of game theory long before it was formally stated in the literature.

Efficiency loss. The price of anarchy concept (without the particular term) was discussed in the computer science literature in Koutsoupias and Papadimitriou [20]; the term “price of anarchy” was introduced in Papadimitriou [34]. In economics, efficiency losses are often referred to as *deadweight losses*, particularly in microeconomic contexts (see the textbook Mas–Colell et al. [24]). A survey of price of anarchy and related efficiency measures can be found Part III of Nisan et al. [31].

4

Externalities

Informally, *externalities* are effects that one agent's actions have on the utility received by another. For example, in the network resource allocation problem discussed in Example 2.4, the rate allocated to one agent increases the latency experienced by the other agents sharing the common resource. There are many other examples of externalities; indeed, interpreted broadly enough, we can view *every* game as exhibiting externalities of some sort, since a single player's actions typically influence the payoffs of other players. These effects can be *negative* (e.g., resource congestion, or the interference one user in a wireless network causes to others) or *positive* (e.g., the performance improvement a user in a peer-to-peer network experiences due to the participation of other users). From our perspective, externalities prove challenging because they often lead to distortions of economic outcomes away from efficiency. In response, there are a range of approaches used to try to manage externalities through appropriate incentives, from pricing to regulation.

In this chapter we discuss externalities in greater detail. We begin in Section 4.1 by defining externalities more precisely and classifying them. In particular, we discuss both positive and negative externalities.

The externalities we discuss in Section 4.1 are primarily based on how the *payoffs* agents receive are affected by each others' actions; in Section 4.2 we focus on settings where externalities arise due to asymmetries in the *information* available to agents.

Given that externalities can lead to inefficient outcomes, we are naturally led to ask: what can be done to mitigate the effect of externalities? In Section 4.3, we discuss what might happen if agents are left to their own devices: in particular, we study the role of *bargaining* in mitigating externalities. This serves as a prelude to a much lengthier discussion in the following chapter on the use of a variety of *incentives* to mitigate the effects of externalities.

4.1 Basic Definitions

As described above, the intuitive notion of an externality is seemingly clear, but also a bit too broad to be economically meaningful. In making the definition of an externality more precise, several subtle issues arise, which we discuss in this section. In addition to defining externalities, we also discuss the classification of externalities as either *positive* or *negative*.

4.1.1 Methodology

Consider a general resource allocation problem of selecting an outcome $\mathbf{x} = (x_1, \dots, x_R)$ for R agents from a set of feasible outcomes \mathcal{X} , where x_r represents agent r 's share of the resource. Loosely, if there are no externalities, then each agent r 's utility for a given allocation depends only on x_r ; i.e., it is a function $U_r(x_r)$, and does not depend on the values of x_s for $s \neq r$. On the other hand, when externalities are present, agents have utilities that depend on both x_r and x_s for $s \neq r$; i.e., the utilities are now of the form $u_r(x_1, \dots, x_R)$.

This is not a precise definition. As noted above, in this sense *every* game exhibits externalities, so the definition does not yield any additional insight. There are more subtle problems as well. For example, note that by changing the labels of each outcome, one can change the dependence of the utilities on these labels. As an extreme example, we could simply label each outcome by setting x_r equal to

the corresponding utility received by each agent; in this case we have $U_r(x_r) = x_r$, and so there is never any dependence on x_s ($s \neq r$) in user r 's utility. A useful definition of externalities should not depend on an arbitrary choice of labels.

A more precise definition of externalities is somewhat subtle. To see this, note that even in the case of allocating a divisible resource as in Example 2.1, one agent's share of the resource reduces the amount of resource available, and so directly has an effect on the utility that can be received by other agents. Is this an externality? If so, then our definition of externality needs to be expanded further, beyond its (already broad) description above. Such an extension would make our definition of externality so broad as to be vacuous.

A more meaningful definition of externality is typically obtained by defining externalities not only in terms of the underlying resource, but within the context of a *market* for that resource. Given a market, an externality is defined as an effect that one agent causes others, *that is not accounted for by the market mechanism*. For example, there are no externalities in a market for the divisible resource in Example 2.1 if a market clearing price is used to reflect the contention for the common capacity constraint. Though this definition is more precise, the term externality is more often used in the looser sense described above (despite the resulting breadth of the definition): to simply reflect the dependence of one agent's utility function on the resources obtained by another, under some "natural" parameterization of the resources.

We conclude by reiterating that externalities may be either *positive* or *negative*. A negative externality refers to a situation where increasing one user's share of a resource *decreases* the utility of other agents; or more broadly, where a higher action by an agent decreases the utility of others. The congestion effect in Example 2.4 is one example of a negative externality. By contrast, a positive externality refers to the situation where increasing one user's resource consumption (or her action) increases the utility of other agents. Peer-to-peer filesharing systems exhibit positive externalities: when more users participate, the system performance typically improves for everyone. Note that in a given setting, both types of externalities may be present; for example,

in a filesharing system, additional users may also increase congestion (a negative externality).

4.1.2 Examples

We give two examples illustrating the different types of externalities (negative and positive).

Example 4.1 (Interference externalities). Consider a wireless network where a set of R users share a common frequency band. Each user corresponds to a unique transmitter/receiver pair. User r sets a power level P_r and transmits over the entire band treating all interference as noise. The performance of each user r is given by a utility $u_r(\gamma_r)$ that is an increasing function of their signal-to-interference plus noise ratio (SINR) γ_r , which is given by

$$\gamma_r = \frac{h_{rr}P_r}{n_0 + \sum_{s \neq r} h_{sr}P_s},$$

where h_{sr} denotes the channel gain from transmitter s to receiver r , and n_0 is the noise power. Suppose that each user can select a power P_r from the interval $[0, P_{\max}]$ to maximize their own utility. Since γ_r is increasing in P_r , it follows that each user would transmit at the maximum power P_{\max} , which can result in a total utility that is much smaller than the maximum achievable utility.

The reason for this outcome is that while increasing an individual's power increases their own SINR, it decreases the SINR of every other user, i.e., this is a negative externality. Each player does not account for this externality in making her own decision; a common expression for this result is that players are not "internalizing the externality." It is instructive to compare this to the congestion externality in Example 2.4. With the congestion externality, all users of the resource experience the same disutility due to the externality, while with this interference model the disutility varies across the users depending on the cross channel gains h_{sr} .

Example 4.2 (Network effects). Consider the following high level abstraction of an online gaming system with R participants. Each participant r exerts an amount of “effort” $e_r \in [0, 1]$; for example, we could interpret this to correspond to the time spent online, with 1 the maximum “effort” possible. Suppose that this effort incurs a cost to the participant of $c_r(e_r)$; for simplicity, we assume $c_r(e_r) = e_r^2$. On the other hand, suppose that the benefit to participant r is proportional to *both* her own effort e_r , as well as the total effort of other participants, $\sum_{s \neq r} e_s$. Thus the total utility of each participant r is

$$U_r(e_r, \mathbf{e}_{-r}) = \alpha_r e_r \sum_{s \neq r} e_s - e_r^2,$$

where $\alpha_r > 0$. In this case if agent r increases her effort, she improves the payoff of each other user s . Thus this is a model with positive externalities.

Provided that $\alpha_r(R - 1) + \sum_{s \neq r} \alpha_s > 2$ for all r , the total utility of all agents is maximized if each agent chooses maximum effort, i.e., $e_r = 1$ for all r . On the other hand, suppose that for some user r , $\alpha_r(R - 1) < 2$. In this case, if all users $s \neq r$ exert effort $e_s = 1$, user r 's marginal change in utility at $e_r = 1$ is $\alpha_r(R - 1) - 2 < 0$. In particular, user r has an incentive to reduce her effort and *free ride* on the effort of the other users. This argument can be refined to show that *no* Nash equilibrium can be Pareto efficient; we leave the details as an exercise for the reader. Indeed, if $\alpha_r(R - 1) < 2$ for *all* r , then the unique equilibrium is for each participant to exert zero effort, resulting in a 100% efficiency loss.

This type of externality is also referred to as a *network externality* or a *network effect*, referring to the fact that the value of the service to a user depends on the number of other users in the “network.” As the name suggests, such effects are a common feature of many networking scenarios ranging from the use of a particular networking technology to social networking applications that operate on a network. Note that in this example if each user exerts her maximum effort, then the value of the network grows proportional to the square of the number of users; this is an observation known as “Metcalfe’s law,” which Bob Metcalfe originally put forward in the context of the value of Ethernet networks.

4.1.3 Discussion

Over-consumption and negative externalities vs. free-riding and positive externalities. Broadly speaking, in the absence of any corrective economic mechanism, negative externalities tend to lead to *over-consumption* of resources, while positive externalities tend to lead to *free-riding* on the efforts of others (and under-provisioning of effort as a result). With negative externalities, self interested users do not account for the negative effects they have on other users, and this leads them to over-consume compared to the efficient allocation. This is the reason for the tragedy of the commons in Example 3.9, and the more extreme failure in the interference example above. On the other hand, with positive externalities, agents receive positive benefits from others' actions, potentially leading to their own inaction. The free-riding effect in the peer-to-peer example is a canonical version of this phenomenon.

4.2 Informational Externalities

In the examples in the preceding section, externalities are modeled in terms of the payoffs assigned to each user. Such effects are sometimes referred to as *payoff externalities*. In this section, we briefly describe another significant form of externality: *informational externalities*, which arise when there are asymmetries in the information available to different players.

4.2.1 Methodology

In this section, we consider a resource allocation problem with incomplete information among the players. As in the previous section, given R agents, the problem is to select an allocation $(x_1, \dots, x_R) \in \mathcal{X}$ (where \mathcal{X} represents the feasible set of allocations). We assume that each agent r 's value of the allocation depends not only on her own share of the allocation x_r , but also on a random variable $\omega \in \Omega$ chosen according to a known probability distribution; i.e., agent r 's payoff is given by $\Pi_r(x_r, \omega)$. The random variable ω is sometimes referred to as "the state of the world." Each agent r receives a signal τ_r about the state of the world, where Θ_r is the set of possible signal values for agent r ; τ_r is a random variable dependent on ω .

If we allow each agent to specify x_r , then we can view this as a Bayesian game where Θ_r is the type space of agent r . With the preceding assumptions, agent r 's payoff does not depend on x_s for $s \neq r$, and so there are no payoff externalities. However, player r 's *a priori* expected payoff *may* depend on the signal of user $s \neq r$. Specifically, this will be the case if

$$E_\omega[\Pi_r(x_r, \omega)|\tau_r] \neq E_\omega[\Pi_r(x_r, \omega)|\tau_r, \tau_s].$$

When this is the case, agent s is said to impose an *informational externality* on agent r . For example, in the formulation of the “market for lemons” as a Bayesian game in Example 3.8, the seller imposed such an externality on the buyer, leading to the potential market failure. As with payoff externalities, the effect of an informational externality can be either positive or negative. Indeed, whether this effect is positive or negative may in fact depend on the specific realization of the state of the world.

4.2.2 Example

If agents do not simultaneously make decisions, then informational externalities can result in one agent's actions influencing the beliefs of other agents — for better or worse. This is illustrated in the following seminal example.

Example 4.3 (Herding). Consider a set of agents, who each have the option of taking one of two irreversible actions $a_r \in \{A, B\}$; for example, the action may correspond to placing a bet on an outcome. The value of each action depends on an unknown state of the world ω that also takes values in the set $\{A, B\}$, which are *a priori* equally likely. If a player chooses an action that matches the true state of the world they receive a utility of 1, otherwise, they receive a disutility of -1 . Each player seeks to maximize their expected payoff.

Without any further information, each player would be indifferent on the choice of action, as the expected payoff would be zero for either. Now suppose that each player also receives a signal, $\tau_r \in \{A, B\}$, which gives her additional knowledge of the state of the world. Specifically, τ_r indicates the true state of the world with probability $1 - \epsilon > 1/2$;

and indicates the opposite state with the remaining probability ϵ . Each player's signal is private information; however, each player is aware that the other players are also receiving this information. Given this private signal, player r is led to believe that the state of the world is equal to τ_r with probability $1 - \epsilon$, so her expected payoff is maximized by choosing $a_r = \tau_r$.

Finally, suppose that the players sequentially decide their actions after observing the action of all players before them. First consider a player 2 that observes player 1 taking action A before deciding on her own action. This informs player 2 that player 1's signal must have been $\tau_1 = A$, which is information that player 2 can use to further revise her belief. In particular, if player 2's signal was $\tau_2 = B$, then player 2's belief that the state of the world is B now becomes

$$\Pr(\Omega = B | \tau_1 = A, \tau_2 = B) = 1/2;$$

thus player 2 becomes indifferent between the two actions. In other words, even though the actual value of the good to player 2 is independent of the action of player 1, player 1's action can affect player 2's belief about this value. Moreover, player 1 does not account for this effect when making her decision.

Even more surprising behavior can arise if we continue. In particular, suppose that when indifferent, a player follows her own signal; and suppose that the first two players choose A . It can then be shown that *every subsequent player will choose A regardless of their own signal!* (We leave the proof of this remarkable fact to the reader; see also the endnotes.) This behavior arises because the first two players choosing A reveals to the next player that they both had signals equal to A . This information overrides the third player's own signal, leading her to choose A as well. The fourth player now has the *same* decision problem as the third player (since the third player's action conveys no new information); so the fourth player (and all subsequent players) will all choose A , regardless of their own signal.

The preceding effect referred to as *herding* behavior or an *information cascade*. Information cascades illustrate a particularly surprising form of information externality: in a cascade, individuals are exerting *negative* informational externalities on their successors. By acting based

on past actions rather than their own signals, each player precludes her successors from benefiting from the information contained in her signal. It is worth noting that herds and cascades can be asymptotically efficient (i.e., everyone eventually taking the correct decision) or inefficient (i.e., everyone eventually taking the incorrect decision).

4.2.3 Discussion

Mitigating information externalities. How are such externalities addressed? One solution to this type of problem is to encourage the agents to reveal their private information before anyone acts. Of course a concern here is that agents may not truthfully reveal their information; the theory of mechanism design (see Chapter 5) provides some techniques to address this issue. In the herding context, one possible simple solution is to have the first N individuals make decisions in isolation *without* observation of the past. This effectively gives successive individuals the opportunity to observe N *independent* signals, significantly increasing the likelihood of taking the efficient action. Such strategies are often pursued in practical settings; decision makers benefit from a diversity of perspectives to avoid “groupthink.”

Long-run interactions and reputation. In many examples with information externalities, players are repeatedly interacting, but against a set of players that changes over time. For example, in online marketplaces such as eBay, buyers and sellers interact routinely, but not necessarily with the same counterparty. In such settings, *reputation systems* can often serve as a “passive” mechanism to align incentives with efficient outcomes despite information asymmetries.

Informally, a reputation system collects and aggregates information about an agent’s previous behavior, and makes this summary available to all system participants. For example, in eBay, the marketplace requests feedback from market participants on each transaction; this gives sellers an incentive to exert effort to improve their reputation. In particular, the seller has an incentive to develop a reputation for selling high quality goods so that buyers become willing to trade — overcoming the “market for lemons” scenario. Alternatively, the threat of lost reputation is often enough to persuade sellers to reveal their private

information on the quality of a good through their advertisement of the product.

There are many challenges in developing a reputation system; we cannot be exhaustive here, but briefly mention a few. First, identity management is needed; if agents can easily change their identity, then they can avoid punishment from bad behavior (though at the expense of losing any accumulated reputation). Also, a process dependent on input from system participants to form reputations can be susceptible to collusion; e.g., a group of agents could falsely inflate their own reputations. Transaction costs are typically necessary to prevent such strategies from being profitable.

This discussion illustrates the role that repeated or *long-run* interactions play in mitigating externalities. As discussed in Section 3.6, repeated games typically possess equilibria that can be much more efficient than the equilibria of the corresponding one-shot game. In that section, we demonstrated such an equilibrium for the prisoner's dilemma, based on the principle of punishment for deviation from a desired efficient outcome. Reputation systems help encourage this type of long-run enforcement in decentralized settings where individuals may not be able to directly keep track of each other's past behavior. We conclude with a warning, however: as we also discussed in Section 3.6, repeated games typically possess *many* equilibria, some of which are highly inefficient; similar issues can arise in the presence of reputation effects.

4.3 Bargaining

In this section, we begin our discussion of how externalities might be mitigated and efficient outcomes recovered. While many approaches to addressing externalities concern active intervention — e.g., explicit incentive mechanisms, markets, etc. — in this section, we consider a simpler alternative: what happens if we simply allow individuals to *bargain* with each other regarding the amount of externality produced? This approach was extensively studied by the economist Ronald Coase in the 1950s, in the context of the allocation of wireless spectrum.

4.3.1 Methodology

To motivate the bargaining approach, we fix a simple scenario. In particular, suppose we must determine a resource allocation (x_1, x_2) to two agents, with $x_1 \in [0, 1]$ and $x_2 \in [0, 1]$. Each user has a utility of the form $U_r(x_1, x_2)$ which is increasing in the agent's own allocation but decreasing in the other agent's allocation; in other words, there is a negative externality. We also assume preferences are quasilinear, so that the utilities are valued in units of a common currency.

Suppose that the unique efficient allocation is $x_1 = 1$ and $x_2 = 0$; however, when each user chooses their own usage, the resulting allocation is $x_1 = 1$ and $x_2 = 1$. Thus we know that $U_2(1, 1) > U_2(1, 0)$, and $U_1(1, 1) + U_2(1, 1) < U_1(1, 0) + U_2(1, 0)$. Rearranging terms we obtain:

$$U_1(1, 0) - U_1(1, 1) > U_2(1, 1) - U_2(1, 0) > 0.$$

Under a bargaining approach, we assume that the two providers can bargain with each other over the amount of externality each imposes on the other. For example, suppose that agent 1 offers agent 2 a payment p such that

$$U_1(1, 0) - U_1(1, 1) > p > U_2(1, 1) - U_2(1, 0)$$

to change its decision from 1 to 0. Agent 2 should be willing to accept this offer as the payment exceeds its loss in utility; and agent 1 should be willing to make such an offer as the payment is less than the utility it gains. After such a transaction, the agents will have arrived at the efficient outcome.

The preceding result is an example of the *Coase theorem*. Informally, the Coase theorem states that if trade for an externality can occur, then in the absence of frictions, bargaining will lead to an efficient outcome. (For a more formal statement of the theorem, see the endnotes.)

A prerequisite for the Coase theorem to apply is the existence of well-defined, verifiable, and enforceable *property rights* for the assets being traded. In other words, agents must be able to clearly value the externality, and ensure that the counterparty follows through on their side of the bargain. Indeed this result is often stated in terms of the allocation of property rights.

Referring to the preceding example, we were assuming that initially each agent had the right to select any allocation in $[0,1]$ for herself, but not the other agent. After accepting the above bargain, agent 2 is essentially relinquishing this right to agent 1. Now suppose we change the initial allocation of rights and give agent 2 the right to select *both* her own allocation as well as agent 1's allocation. In this case, without bargaining agent 2 would choose $x_1 = 0$ and $x_2 = 1$. Suppose also that $(1,0)$ is more efficient than $(0,1)$; in that case, we obtain:

$$U_1(1,0) - U_1(0,1) > U_2(0,1) - U_2(1,0).$$

Therefore agent 1 can again offer agent 2 an acceptable price between these two limits to obtain the right to set x_1 to 1, while decreasing x_2 to zero. The point of this example is that the initial allocation of property rights does not matter provided that they can be traded without frictions; this is an alternative way of stating the Coase theorem.

4.3.2 Discussion

Frictions and bargaining inefficiency. The Coase theorem may seem to suggest a very simple market structure for well defined assets: simply allocate the assets arbitrarily with well defined property rights, and let the agents bargain. The important point that such a conclusion overlooks is the presence of *frictions* that can impede efficient bargaining. For example, there may be costs and delays involved in bargaining. These delays may be due to the time it takes to agree on a bargain, as well as the search time and effort needed to find a counterparty to bargain with. Further, such frictions tend to increase with the number of parties involved in bargaining. For example, in the case of Internet peering relationships or wireless interference, the externalities are multilateral and may require bargaining between multiple agents to resolve efficiently. Such multilateral bargaining is typically costly to effectively execute, and consequently efficiency may be lost.

More importantly, we emphasize that the Coase theorem does not specify a *protocol* for bargaining. In this sense, the Coase theorem is more of an “existence” proof: it argues that *some* protocol exists for resolving *ex ante* inefficiencies. In practice, a significant amount of effort

is invested in constructing mechanisms that resolve *ex ante* inefficiencies; indeed, most of market design is concerned with precisely this goal in a multilateral setting. A central concern of such mechanisms is typically to reduce the frictions present in any given resource allocation setting, so that individuals can reasonably “discover” the efficient allocation *ex post*.

Asymmetric information. Another friction in bargaining is due to incomplete or asymmetric information, i.e., situations where agents do not exactly know each other’s valuations. Returning to the previous example, suppose that $U_1(1,0) - U_1(1,1) = 3$ and $U_2(1,1) - U_2(1,0) = 2$; thus agent 1 should be willing to pay some price p slightly more than 2, and agent 2 should accept such a price. Now suppose instead that agent 1 only knows that $U_2(1,1) - U_2(1,0)$ is distributed uniformly on the interval $[0,3]$; in this case, agent 1’s expected benefit is maximized if she makes agent 2 an offer of 1.5 — but this offer will not be accepted, even though the overall welfare would improve if trade occurred.

This example is similar in spirit to the failure to obtain efficient trade in the “market for lemons” of Section 2.6. Of course, this is just one possible way the two agents can bargain; however, a result due to the economists Myerson and Satterthwaite shows that the underlying problem is fundamental: namely in the presence of incomplete information there is no bargaining procedure that can always guarantee the efficient outcome (see the endnotes). This discussion suggests that simply relying on bargaining alone may not be sufficient to achieve efficient outcomes.

4.4 Endnotes

Basic definitions. Externalities are covered in most standard texts on microeconomic theory including those listed in the endnotes to Chapter 3.

Informational externalities. Models of herding due to informational externalities originally appear in the work of Banerjee [4] and Bikhchandani et al. [6].

A more thorough overview of reputation systems can be found in Chapter 27 of Nisan et al. [31].

Bargaining. The idea of bargaining over externalities is attributed to Coase [9]. Myerson and Satterthwaite's work on asymmetric information in bargaining appeared in Myerson and Satterthwaite [29].

5

Incentives

The preceding chapter highlighted the distortions away from efficient outcomes that often result in economic systems with externalities. In this chapter, we take a *design* point of view in dealing with externalities: in particular, we ask what *incentives* we can provide to agents to encourage them to align their own behavior with efficient outcomes. Proper design of incentives can mitigate or entirely eliminate the distortions in economic outcomes that arise due to externalities.

Broadly, the design of incentives that lead individuals to align their behavior with desired outcomes is referred to as *mechanism design*. In a resource allocation setting, a “mechanism” refers to the rules of the allocation procedure including a specification of what information the agents send to the mediator, how the mediator allocates the resource based on this information, and what payments if any are required from the agents for their allocation. In essence, such a mechanism is specifying the strategy spaces and payoff functions in a game which the agents will play; for this reason it is sometimes useful to think of mechanism design as “inverse” game theory. Common examples of such mechanisms are various types of price mechanisms (e.g., as described in Section 2.5) and auctions (such as the first-price auction discussed in Example 3.7).

In this chapter, we discuss a range of approaches to designing incentives. We start with the bluntest instrument available to the mechanism designer: regulation of behavior (Section 5.1). Next, we discuss a somewhat more sophisticated approach to providing incentives, inspired by the discussion of prices in Section 2.5. In particular, we discuss in detail *Pigovian taxes* that can lead individuals to “internalize” the externality they impose on others (Section 5.2).

Both Pigovian taxes and regulation suffer from the pitfall that they may depend on prior information about the utilities of participants, and we may not possess this information when we design the incentive scheme. To address this issue, in Section 5.3 we describe the celebrated *Vickrey–Clarke–Groves* (VCG) mechanisms; this class of mechanisms demonstrates that by properly structuring the payment mechanism, it is possible to align individual incentives with overall welfare maximization, even when the preferences of agents are unknown.

The last two sections of the chapter focus on two specific settings of interest: auctions and principal-agent problems. We discuss some basics of auction design in Section 5.4. We conclude with a brief discussion of contract theory and the principal-agent problem in Section 5.5.

5.1 Regulation

We first consider a simple, “brute force” approach to dealing with externalities: *regulate* the behavior of agents to control the level of distortion away from the efficient solution. Such regulation may be done by the government; e.g., the Federal Communications Commission (FCC) in the United States regulates the behavior of wireless devices in part to reduce negative externalities due to interference. However, regulations need not be imposed only by government; a resource provider can also impose rules that users of the system need to follow. For example, when Internet service providers place caps on data usage, this is done in part to limit the congestion externality due to heavy data users.

5.1.1 Methodology

The simplest form of regulation to simply place a limit on the externality-producing activity of an agent. In the case of a negative

externality, this should be an upper bound, while for positive externalities this should be a lower bound. For example, in the case of interference, one can place a limit on the transmission power of each agent, while in the example of positive externalities in Example 4.2 we might place a lower bound on the effort contributed by each agent.

To see the effect of such a limit, we consider a generic resource allocation setting of the following form. A resource allocation $\mathbf{x} = (x_1, \dots, x_R)$ is to be selected for a set of R agents, where $x_r \geq 0$ represents the amount of resource allocated to user r . For a given allocation, each agent r receives a payoff given by $\Pi_r(\mathbf{x})$, which depends on the allocation received by every agent. We also assume that agents have quasilinear preferences, so that $\Pi_r(\mathbf{x})$ is measured in a common unit of currency across all users. For each r , assume that $\Pi_r(\mathbf{x})$ is concave in \mathbf{x} and that as $x_r \rightarrow \infty$, $\Pi_r(\mathbf{x}) \rightarrow -\infty$ so that at an efficient allocation each agent consumes a finite amount of resource. For example, the model in Example 2.4 fits into the framework, where $\Pi_r(\mathbf{x})$ is the difference between the utility an agent derives from its allocation and the latency cost.

Suppose in the preceding setting that the externality is *negative*. Let \mathbf{x}^e denote the equilibrium allocation without any regulatory constraints. Now consider adding a constraint $x_r \leq \bar{x}_r$ for each agent r , and then allowing each agent to maximize her own payoff given this constraint, i.e., over $x_r \in [0, \bar{x}_r]$. Clearly, if $\bar{x}_r > x_r^e$ for all r , these new constraints have no effect and agents choose the unconstrained equilibrium allocation. However, from the assumed concavity of the payoff it follows that if $\bar{x}_r < x_r^e$ for all r , then there will now be a Nash equilibrium in which each agent selects $x_r = \bar{x}_r$. For example, if $\bar{x}_r = x_r^*$, this would result in the efficient allocation — though this would require the regulator to know the utilities beforehand, which is often too strong an assumption.

5.1.2 Discussion

Knowledge of preferences. As noted previously, determining the optimal regulation will typically require knowledge of each user's utility function as well as other parameters of the system. Such knowledge

is typically not available, in which case the regulation may not have the desired effect. Thus effective regulation brings with it a significant informational hurdle of acquiring preference information in the first place.

Recall that mechanism design is concerned with the design of economic environments that lead to efficient outcomes; through this lens, regulation may be viewed as a very blunt form of mechanism design. A key consideration in mechanism design is that the rules of the allocation procedure, and the actions available to participants, are all typically designed *before* the mechanism designer has complete knowledge of who the participants will be, or what their preferences might be. Thus there is typically a tension in mechanism design between the individual participants' incentives, and the overall desired outcome of the mechanism designer.

This challenge is often addressed by having users report their utility functions, and then charging them a price so that they are incentivized to do this truthfully; see, e.g., Section 5.3. Not surprisingly, there is a significant cost to be borne for this additional power: such mechanisms rapidly become complex, and greatly increase both the amount of information the resource manager must collect, as well as the computational burden they must bear to run the mechanism. Though regulation is a coarse solution to mitigate externalities, it has the advantage of being much simpler than most other mechanism design approaches.

Regulations and system engineering. For engineered systems, some externality-limiting “regulations” are typically enforced as part of the system design. To illustrate this point, let us consider allocating a band of wireless spectrum to a set of transmitter/receiver pairs as in Example 4.1. Instead of sharing the band by having each user spread their signal over the entire band, we could instead subdivide the band into orthogonal frequency bands (or assign users orthogonal time slots). With such a change, there is no longer an interference externality; rather, each user affects the others only through their competition for individual frequency bands. The spectrum can now be viewed as a divisible resource; a single price can be used to price the negative externality for this shared resource (similar to Section 2.5), ensuring no externalities remain in the market.

Although this approach leads to simpler market design, from an engineering perspective it may not be ideal. In particular, if the transmitter/receiver pairs are far enough apart, then the earlier power sharing model will better exploit frequency reuse and result in better spectrum utilization. Other approaches could also be used such as allowing the reuse of spectrum when users are “sufficiently” far apart, where the reuse distance is again a way to regulate the externalities imposed on the system. This discussion highlights the importance of understanding whether the rigid limitations imposed by regulations are worth the potential efficiency losses they incur.

5.2 Pigovian Taxes

In this section, we begin discussing more tailored approaches to mitigate the distortions away from efficient outcomes due to externalities. For a congestion externality, we have already alluded to one such solution in Section 2.5, namely the use of a price or *Pigovian tax* to “internalize the externality.” We now consider this type of approach in more detail.

5.2.1 Methodology

For concreteness, consider the following variation of the generic resource allocation problem presented in the previous section, where $\Pi_r(\mathbf{x})$ is the payoff to agent r at allocation \mathbf{x} . Suppose each agent r selects her share of x_r to maximize her own payoff. This results in a game, which from the assumed convexity has a Nash equilibrium. At any such equilibrium \mathbf{x}^e , agent r ’s allocation must satisfy the following first-order optimality condition:

$$\frac{\partial \Pi_r}{\partial x_r}(\mathbf{x}^e) \leq 0, \quad (5.1)$$

with equality if $x_r^e > 0$.

Let us contrast this with an efficient allocation assuming that the payoffs are measured in a common monetary unit; in other words, agents’ overall utility is quasilinear. The efficient allocation is given

by solving:

$$\begin{aligned} & \text{maximize } \sum_r \Pi_r(\mathbf{x}) \\ & \text{subject to: } \mathbf{x} \geq 0. \end{aligned}$$

From the assumed convexity, a necessary and sufficient condition for \mathbf{x}^* to be efficient is that for each r it satisfy:

$$\sum_{s=1}^R \frac{\partial \Pi_s}{\partial x_r}(\mathbf{x}^*) \leq 0, \quad (5.2)$$

with equality if $x_r^* > 0$.

Comparing these two, suppose that we have an interior equilibrium solution $\mathbf{x}^e > \mathbf{0}$ so that (5.1) holds with equality. Substituting these values into (5.2), it follows that \mathbf{x}^e is efficient if and only if for each r ,

$$\sum_{s \neq r} \frac{\partial \Pi_s}{\partial x_r}(\mathbf{x}^e) = 0.$$

If Π_s did not depend on x_r for $r \neq s$, then this relation would hold; however, when externalities are present these terms are not zero in general, and thus \mathbf{x}^e is not efficient in general. Note that with positive externalities these terms are positive, and for negative externalities they are negative.

Suppose now that each user r is charged a price per unit resource t_r . We refer to such a price as a *Pigovian tax* if charging this price to user r results in her choosing the efficient allocation x_r^* . In other words, when user r maximizes $\Pi_r(\mathbf{x}) - t_r x_r$, she should choose x_r^* . (Note that in the literature, the term “Pigovian tax” may be used somewhat more loosely than this precise notion, to refer to any price mechanism that aims to redirect individual behavior towards more efficient outcomes.)

From the previous discussion, it is straightforward to check that such prices have the form given in the following theorem.

Theorem 5.1 The set of Pigovian taxes t_1, \dots, t_R that support the efficient allocation \mathbf{x}^* are given by:

$$t_r = - \sum_{s \neq r} \frac{\partial \Pi_s}{\partial x_r}(\mathbf{x}^*).$$

Note that with positive externalities, t_r is negative in which case these “taxes” are actually “subsidies.”

5.2.2 Examples

Three examples of pricing externalities follow.

Example 5.1 (Tragedy of the commons avoided). Consider the model in Example 3.9 for sharing a link with capacity C among R users. Each user receives a utility of $U_r(x_r) = x_r$ from their share and experiences a latency given by $\ell(y) = y/C$. When the users simply optimize their own payoff $\Pi_r(\mathbf{x}) = U_r(x_r) - \ell(y)$, we saw in Example 3.9 that the resulting equilibria exhibit an efficiency loss that approaches 100% as R increases.

The Pigovian tax for each user r in this example is given by

$$\begin{aligned} t_r &= - \sum_{s \neq r} \frac{\partial \Pi_s}{\partial x_r}(\mathbf{x}^*) \\ &= (R - 1) \frac{d\ell}{dy}(x_1^* + \dots + x_R^*) \\ &= \frac{(R - 1)}{C}. \end{aligned}$$

Hence, if each user is charged a price of $(R - 1)/C$ per unit resource, they internalize the externality, resulting in the efficient allocation.

Note in this example, the Pigovian tax t_r is the same for each user r , i.e., a uniform price per unit resource is sufficient. This arises from the symmetry of the externality: each user incurs the same marginal disutility given an increase in the rate of any user. Further, the price does not depend on the equilibrium allocation \mathbf{x}^* , which is due to the assumed linearity of the latency function. As illustrated by the next example, in general neither of these properties hold.

Example 5.2 (Interference prices). Consider the model for interference externalities in Example 4.1. The Pigovian tax for each user r

in this example is given by

$$t_r = - \sum_{s \neq r} \frac{\partial U_s}{\partial P_r}(\gamma_r^*),$$

where γ_r^* is the SINR for user r at the efficient allocation. For each user r , let $I_r = \sum_{s \neq r} h_{sr} P_r$ be the total interference received at user r . Using this, t_r can be written as

$$\begin{aligned} t_r &= - \sum_{s \neq r} \frac{\partial U_s}{\partial I_s} \frac{\partial I_s}{\partial P_r}(\mathbf{P}^*) \\ &= - \sum_{s \neq r} \frac{\partial U_s}{\partial I_s}(\mathbf{P}^*) h_{rs}. \end{aligned}$$

One can interpret $\pi_s = -\frac{\partial U_s}{\partial I_s}(\mathbf{P}^*)$ as an *interference price* for user s , which indicates that user's marginal decrease in utility due to an increase in the total interference at that user. Given the set of interference prices, the Pigovian tax charged to user r is then given by multiplying the interference price for each user s by the amount of interference user r causes user s (i.e., $h_{rs} P_r$).

Compared to the previous example, here in general, there is a different Pigovian tax charged per unit power to each user. Specifically, at each user the interference prices are weighted differently depending on the channel gains, and the interference prices themselves also vary depending on each user's utility and SINR. This reflects the fact that in this case the underlying externality is not symmetric across the users.

Example 5.3 (Network effects and critical mass). Recall that for the on-line gaming system in Example 4.2, if for each user $\alpha_r(R - 1) < 2$, then the unique equilibrium is for every user to exert zero effort, while if $\alpha_r(R - 1) + \sum_{s \neq r} \alpha_s > 2$ then the total utility is maximized if every user exerts the maximum effort of 1. A Pigovian tax for this problem would be to offer each user an incentive of $\sum_{s \neq r} \alpha_s$ per unit effort, which would result in everyone exerting maximum effort.

In fact offering any incentive t_r such that $\alpha_r(R - 1) + t_r > 2$ would suffice to incentivize users to exert the maximum effort. Note that the

needed incentive decreases with the number of users R ; if α_r is bounded away from 0, for large enough R no incentive is needed at all. This illustrates a common phenomenon of systems with network effects: it is often advantageous to subsidize “early adopters” so as to build a critical mass of customers, after which using the system becomes valuable enough that subsidies are no longer needed.

5.2.3 Discussion

Payments. A careful reader might raise the following objection to the preceding discussion: though Pigovian taxes enable us to achieve the efficient allocation for the original problem, we have changed each user’s net utility by introducing the payment term. If economic models are only being used semantically to guide engineering design, this is not a problem, since we only care about engineering a desired outcome.

However in an economic model, the equilibrium net utility of an agent after paying a Pigovian tax may in fact be *smaller* than the agent’s equilibrium net utility when there are no prices. Is this, then, necessarily a Pareto improvement? To resolve the dilemma, we need to recall why the sum of utilities is the efficiency benchmark in a quasilinear setting (per Section 2.4). If the system moves from an equilibrium outcome to one where total utility is maximized, then there is a way to structure redistributions of wealth to agents so that *everyone* is left better off. In this view, the resource allocator is simply a bookkeeper: if the resource allocator promises a fixed payment (or fixed fee) to each user, where in equilibrium the total amount paid out equals the total payments collected, then we obtain the same total utility as in the efficient solution.

This interpretation typically leads us to ignore any payments made to or from the resource allocator when we consider Pigovian taxes. (See also the discussion of budget balance in the following section.) However, note that in practice, the resource manager (or market operator) typically does not rebate net proceeds in this way. As a consequence, individuals may forecast a potential reduction in their own net utility, and choose not participate. Ensuring participation is a common goal in managing externalities; see, e.g., Section 5.5 for a description of this constraint in the context of contracting.

Price discovery. As illustrated in Example 5.2, in general the Pigovian tax for a problem depends on the efficient allocation, which in turn depends on parameters of the problem such as the utilities of each user. Of course, the key issue in most places where economic models are applied is that this information is not known by the “resource manager,” which then begs the question of how such a price is set in the first place. Here we discuss a few resolutions to this quandary.

One possibility is that Pigovian taxes are obtained via some type of price adjustment process. Indeed, the congestion externality in Example 5.1 can be viewed as a special case of the formulation in Section 2.5. As discussed there, one such pricing mechanism is to adjust the price given the current flows to equal the marginal cost. Such a mechanism does not require the resource manager to know each user’s utility, $U_r(x_r)$, but does require that they have knowledge of the latency function $\ell(x)$. Given this knowledge the resource manager can set the price equal to the marginal change in latency; this only requires measurement of the total flow through the resource. From an engineering perspective, this leads to a natural distributed algorithm for determining the Pigovian tax.

While reasonable for pure engineering design, there are two issues with this interpretation from an economic perspective. First, if one is truly modeling economic agents, then agents may be able to *forecast* how prices will change, and alter their behavior to affect the outcome of the price setting process. Second, we are implicitly assuming all users value latency in the same (known) way, which may not be true. In general, therefore, a more sophisticated approach to managing users’ incentives may be needed.

For other types of externalities, price discovery may be more complicated. For example, in Example 5.2, a resource manager would need to determine not just one price but R interference prices; and each of these prices depend in turn on the local interference conditions at each node, which may be difficult or impossible for a resource manager to measure. An alternative in such a setting is for each agent to determine their *own* interference prices, by setting the price equal to the marginal change in their utility due to interference; the nodes can then announce these prices to other users. Given the prices, if each user updates their transmission powers in a greedy fashion, under appropriate assumptions

it can be shown that the resulting process converges to the Pigovian taxes for this problem. Again this provides an engineering algorithm for allocating resources; however, from an economic point of view it is somewhat problematic. The issue here is that users are also *price setters* and would always have an incentive to announce inflated interference prices, since this would reduce the interference they see.

Monitoring. In some settings actually monitoring agents' activity may be difficult, complicating the use of Pigovian taxes. For example, in the case of interference, it might not be possible for a resource allocator to accurately measure the interference received at each receiver; moreover, given the total interference, it can be difficult to determine each user's contribution to this total. Of course, agents could be asked to report information such as their transmit power or received SINR to help facilitate attribution, but they may not have the incentive to report this information correctly (again raising a mechanism design problem).

5.3 The Vickrey-Clarke-Groves Mechanisms

Pigovian taxes are an appealingly simple solution to management of externalities, but are not without their flaws. Our discussion of price discovery and monitoring in the preceding section highlights one significant challenge in particular: achieving efficient outcomes requires knowledge of preferences. If agents are strategic, they may distort the equilibrium outcome away from efficiency; Pigovian taxes may not be enough to preclude such an incentive, e.g., in the case where users are also price setters.

In this section, we describe the most famous approach in mechanism design for resolving this issue, captured by the *Vickrey-Clarke-Groves* (VCG) class of mechanisms. These are mechanisms that guarantee efficient allocation of resources in environments where utilities are quasilinear (cf. Section 2.4). The key insight behind VCG mechanisms is that by structuring payment rules correctly, individuals can be incentivized to truthfully declare their utility functions to the market. The mechanism can then efficiently allocate resources with respect to these (truthfully) declared utility functions. As we discuss, VCG

mechanisms may be viewed as a powerful extension of the Pigovian taxation principle derived in the preceding section.

5.3.1 Methodology

To illustrate the principle behind VCG mechanisms, we turn to a variation of Example 3.2, where we are allocating a single resource among R users. In this case, we do not consider any latency function, but consider a simpler model where the resource simply has a fixed capacity C that is to be allocated among competing users.

Each user r has quasilinear utility; in particular, if the allocated amount is x_r and the payment to user r is t_r , then her utility is:

$$U_r(x_r) + t_r.$$

We assume that for each r , over the domain $x_r \geq 0$ the valuation function $U_r(x_r)$ is concave, strictly increasing, and continuous; and over the domain $x_r > 0$, $U_r(x_r)$ is continuously differentiable. Furthermore, the right directional derivative at 0, denoted $U_r'(0)$, is finite. We let \mathcal{U} denote the set of all functions satisfying these conditions.

Given complete knowledge and centralized control of the system, the welfare maximizing (and Pareto efficient) allocation is given by the following optimization problem (cf. Section 2.4):

$$\text{maximize } \sum_r U_r(x_r) \quad (5.3)$$

$$\text{subject to } \sum_r x_r \leq C; \quad (5.4)$$

$$\mathbf{x} \geq 0. \quad (5.5)$$

Since the objective function is continuous and the feasible region is compact, an optimal solution $\mathbf{x} = (x_1, \dots, x_R)$ exists. If the functions U_r are strictly concave, then the optimal solution is unique, since the feasible region is convex. The main question we now want to address is how a manager should allocate these resource when it does not have *a priori* knowledge of the users' utility functions.

The basic approach in a VCG mechanism is to let the strategy space of each user r be the set \mathcal{U} of possible valuation functions, and make a

payment t_r to user r so that her net payoff has the same form as the social objective (5.3). Thus, in a VCG mechanism, each user is simply asked to *declare* their valuation function; of course, if the payments t_r are not structured properly, there is no guarantee that individuals will make truthful declarations. For each r , we use \hat{U}_r to denote the declared valuation function of user r , and use $\hat{\mathbf{U}} = (\hat{U}_1, \dots, \hat{U}_R)$ to denote the vector of declared valuations.

If user r receives an allocation x_r and a payment t_r , the payoff to user r is:

$$U_r(x_r) + t_r.$$

On the other hand, the social objective (5.3) can be written as:

$$U_r(x_r) + \sum_{s \neq r} U_s(x_s).$$

Comparing the preceding two expressions, the most natural means to align user objectives with the social planner's objectives is to *define the payment t_r as the sum of the valuations of all users other than r* .

Formally, given a vector of declared valuation functions $\hat{\mathbf{U}}$, a VCG mechanism chooses the allocation $\mathbf{x}(\hat{\mathbf{U}})$ as an optimal solution to (5.3)–(5.5) given $\hat{\mathbf{U}}$, i.e.:

$$\mathbf{x}(\hat{\mathbf{U}}) \in \arg \max_{\mathbf{x} \geq 0: \sum_r x_r \leq C} \sum_r \hat{U}_r(x_r). \quad (5.6)$$

The payments are then structured so that:

$$t_r(\hat{\mathbf{U}}) = \sum_{s \neq r} \hat{U}_s(x_s(\hat{\mathbf{U}})) + h_r(\hat{\mathbf{U}}_{-r}). \quad (5.7)$$

Here h_r is an arbitrary function of the declared valuation functions of users other than r ; since user r cannot affect this term through the choice of \hat{U}_r , she chooses \hat{U}_r to maximize:

$$P_r(x_r(\hat{\mathbf{U}}), t_r(\hat{\mathbf{U}})) = U_r(x_r(\hat{\mathbf{U}})) + \sum_{s \neq r} \hat{U}_s(x_s(\hat{\mathbf{U}})).$$

Now note that given $\hat{\mathbf{U}}_{-r}$, the above expression is bounded above by:

$$\max_{\mathbf{x} \in \mathcal{X}} \left[U_r(x_r) + \sum_{s \neq r} \hat{U}_s(x_s) \right].$$

But since $\mathbf{x}(\hat{U})$ satisfies (5.6), user r can achieve the preceding maximum by truthfully declaring $\hat{U}_r = U_r$. Since this optimal strategy does not depend on the valuation functions ($\hat{U}_s, s \neq r$) declared by the other users, we recover the important fact that in a VCG mechanism, truthful declaration is a weak dominant strategy for user r . This discussion is summarized in the following proposition.

Proposition 5.2. Consider a VCG mechanism defined according to (5.6) and (5.7). Then, declaring $\hat{U}_r = U_r$ is a weak dominant strategy for each user r . Furthermore, under these strategies, the resulting allocation is efficient.

For our purposes, the interesting feature of the VCG mechanism is that it elicits the true utilities from the users, and in turn (because of the definition of $\mathbf{x}(\hat{U})$) chooses an efficient allocation. The feature that truthfulness is a dominant strategy is known as *incentive compatibility*: the individual incentives of users are aligned, or “compatible,” with overall efficiency of the system. The VCG mechanism achieves this by effectively paying each agent to tell the truth. The significance of the approach is that this payment can be properly structured even if the resource manager does not have prior knowledge of the true valuation functions.

An important special case of the VCG mechanism arises when the function h_r is defined as:

$$h_r(\hat{U}_{-r}) = - \sum_{s \neq r} \hat{U}_s(x_s(\hat{U}_{-r})),$$

where by an abuse of notation we use $\mathbf{x}(\hat{U}_{-r})$ to denote the optimal allocation of the resource if user r were not present in the system at all. Observe that in this case the VCG mechanism is *as if* user r makes a payment *to* the resource manager, in the amount:

$$-h_r(\hat{U}_{-r}) - t_r(\hat{U}) = \sum_{s \neq r} \left(\hat{U}_s(x_s(\hat{U}_{-r})) - \hat{U}_s(x_s(\hat{U})) \right). \quad (5.8)$$

Note the right-hand side is always nonnegative. Further, user r makes a positive payment if and only if she is *pivotal*, that is, if her presence alters the resource allocation among the other players; for this

reason this particular form of VCG mechanism is sometimes called the *pivot* mechanism or the *Clarke pivot rule*. This formulation of the VCG mechanism is quite intuitive: the payment directly measures the *negative externality* that agent r imposes on other agents (according to their reported valuations), and collects a fee to compensate those agents accordingly. In this sense, the payment in the pivot mechanism serves the same purpose as a Pigovian tax in the preceding section, and can be viewed as a generalization of that concept. Indeed, VCG payments force agents to “internalize the externality,” exactly as Pigovian taxes do. However, while a Pigovian tax indicates the “marginal externality” an agent imposes on other agents, the VCG payment indicates the “total externality” imposed.

Though we have defined this mechanism for allocating a single shared resource as in (5.3)–(5.5), this approach directly generalizes to any resource allocation problem in quasilinear environments, where the goal is to maximize the utilitarian social welfare. More precisely, let \mathcal{X} denote the set of all possible ways to allocate a given set of resources and let $U_r(\mathbf{x})$ denote user r 's valuation for an allocation $\mathbf{x} \in \mathcal{X}$. Then a VCG mechanism for allocating these resources is given by asking each agent to declare a valuation function \hat{U}_r , which in principle could depend on the entire allocation \mathbf{x} . Given the declared valuations, the allocation rule in (5.6) becomes

$$\mathbf{x}(\hat{U}) \in \arg \max_{\mathbf{x} \in \mathcal{X}} \sum_r \hat{U}_r(\mathbf{x}).$$

and the payment rule in (5.7) is modified in the same way.

5.3.2 Examples

Here we present two examples of the VCG mechanism.

Example 5.4 (VCG with interference externalities). We first consider applying the VCG mechanism in a setting with R interfering wireless agents as in Example 4.1. In this case, \mathcal{X} is the set of all power allocations P_1, \dots, P_R that satisfy $0 \leq P_r \leq P_{\max}$ for each r . Each user r receives a utility $U_r(\gamma_r)$ that depends on their received SINR γ_r , which depends in turn on the vector of powers $\mathbf{P} = (P_1, \dots, P_R)$. To

implement a VCG mechanism, the resource allocator needs each user to specify her utility as a function of the entire power profile \mathbf{P} . Let V_r be this report from user r . Given these reported utilities, a VCG mechanism determines the power allocation by solving:

$$\begin{aligned} & \text{maximize } \sum_r V_r(P_1, \dots, P_R) \\ & \text{subject to: } 0 \leq P_r \leq P_{\max}, \quad \text{for all } r. \end{aligned}$$

To determine the VCG payment for user r in (5.8) requires also solving the problem

$$\begin{aligned} & \text{maximize } \sum_{s \neq r} V_s(P_1, \dots, P_R) \\ & \text{subject to: } P_r = 0, \\ & \quad 0 \leq P_s \leq P_{\max}, \quad \text{for all } s \neq r, \end{aligned}$$

which determines the power allocation each user s receives if agent r is absent. The price charged to agent r in (5.8) is then given by the difference between the sum utility the other agents receive under this power profile, and the sum utility they receive under the allocated powers.

Though such a mechanism is incentive compatible and efficient, note that it requires solving $R + 1$ optimization problems, each of which in general will not be convex. It also requires each agent to report a function of R variables, which depends on parameters such as channel gains among the agents that can be difficult to measure.

Example 5.5 *Second price auction.* In this example, we consider allocation of an *indivisible* good. We consider a setting where we have a single copy of a good to allocate among one of R agents. In this case, \mathcal{X} is a discrete set consisting of the R possible ways of allocating the resource. Suppose that the agent who obtains the good earns a valuation of u_r , while all other agents receive zero value. In this case, the VCG mechanism asks each player to simply report the single number u_r . The allocation rule reduces to giving the capacity to an agent

with the largest declared valuation. The payment in (5.8) is equal to the second highest declared valuation.

It is easier to describe this in more practical terms. In the resulting VCG mechanism, each agent r is asked to declare a valuation; we refer to this declaration as the *bid* of user r . The rule above then reduces to the following mechanism: the good is allocated to the agent who makes the highest bid, and this agent pays the *second* highest bid. This is a common auction format known as the *second price auction* (or Vickrey auction), and our analysis leads to the celebrated result that truthfully declaring one's value is a dominant strategy in the second price auction.

5.3.3 Discussion

At first glance, it appears that the VCG mechanism provides an elegant solution to any resource allocation problem, as it encourages agents to correctly reveal their utility functions, and in turn yields fully efficient allocations. Indeed, it can be shown that any mechanism with these properties is in a certain sense equivalent to VCG. More precisely, any efficient, incentive compatible mechanism must result in payments that are structurally the same as those charged by a VCG mechanism. At the same time, there are several reasons why VCG mechanisms may not be entirely desirable; we describe three here.

Budget balance. As in our discussion of externalities in the preceding section, efficiency of the VCG outcome implicitly assumes that monetary transfers can be “ignored.” However, if we sum the VCG payments, there is always either a withdrawal or injection of currency into the system. Some other specification of the model must account for the fate of this currency: either a form of rebating to or taxation on the users in the system; or collection of the excess or injection of the shortfall by a disinterested resource manager.

The fact that the VCG payments do not sum to zero is known as a failure of *budget balance*. The Myerson-Satterthwaite theorem (discussed in Section 4.3) can be reinterpreted as showing that in general resource allocation settings, there is no efficient, budget balanced, and incentive compatible mechanism, in which users will voluntarily choose to participate; see the endnotes to Chapter 4 for

details. Circumventing this result typically requires focusing on a more specific class of allocation problems.

Complexity. The VCG mechanism requires agents to submit entire utility functions. In an engineering context, such as rate allocation in networks, it may be considered prohibitive to submit entire utilities. Indeed, in general individuals may not even explicitly know their own utility functions. Further complications arise because it may be difficult for agents to computationally reason about the impact that their declaration has on their eventual allocation. Additionally, in some markets it may even be computationally difficult for the resource manager to solve the optimization problems required to implement the VCG mechanism, e.g., when these involve combinatorial constraints over many goods. As a result, mechanisms where the strategy space is simpler are desirable. Indeed, most market mechanisms observed in practice do not involve submission of complete utilities.

Nonuniform pricing. The VCG mechanism exhibits an intrinsic “nonuniformity” in the nature of pricing. Consider the pivot mechanism described above as an example. It is clear that the per unit price paid by one agent may not be the same as the per unit price paid by another agent, simply because they may differ in the impact their presence has on the overall resource allocation. This is an unusual feature not typically observed in resource markets, where agents see a common unit price for the good. However, the pivot mechanism identifies these prices based on the negative externality an agent imposes on others, and this externality may of course be nonuniform.

5.4 Auctions

A significant subset of the literature on mechanism design and market design focuses on *auctions*. Auctions are one of the oldest and most basic market mechanisms. We consider auction settings here where one (indivisible) item is up for sale, and bidders submit bids (simultaneously) for the item. We focus on developing both the basic insights of the theory, as well as some of the key insights from a practical mechanism design perspective.

5.4.1 Methodology

An auction is a basic format for sale of an item. We consider auctions with R bidders, where bidder r has a valuation $v_r \geq 0$ for the item. The setting we consider is known as a *private values* setting: each bidder has their own valuation for the item, and their utility is directly determined by only this valuation. (We later contrast this with a *common value* setting.)

Each bidder chooses a bid $b_r \geq 0$, giving rise to a composite bid vector $\mathbf{b} = (b_1, \dots, b_R)$. In the model we consider, an auction is specified by two components:

- (1) *The allocation rule.* The first component of the auction determines the winner (if any). In particular, let $x_r(\mathbf{b}) = 1$ if bidder r wins the item, and zero otherwise. Since only one item is for sale, we have the constraint that $\sum_r x_r(\mathbf{b}) \leq 1$. In most of the auctions we consider, the highest bidder will win (if unique); in this case, $x_r(\mathbf{b}) = 1$ if $b_r > b_s$ for all $s \neq r$. (For completeness, such an allocation rule must also specify how to break ties in case multiple agents submit the highest bid.)
- (2) *The payment rule.* The second component of the auction determines how much participants must pay. In particular, let $w_r(\mathbf{b})$ be the total payment that bidder r is asked to make, when the bids are \mathbf{b} .

Bidder's utilities are quasilinear, so in particular when the bid vector is \mathbf{b} , bidder r 's utility is:

$$v_r x_r(\mathbf{b}) - w_r(\mathbf{b}).$$

The mechanism designer's problem is to *design* the allocation and payment rules to achieve a desired outcome in equilibrium (e.g., allocation of the good to the highest valuation bidder, or revenue maximization). Typically, one of three notions of equilibria are studied: weak dominant strategies; Nash equilibrium; or Bayesian–Nash equilibrium.

5.4.2 Examples

In this section, we describe three examples of auctions that can be studied using this framework.

Example 5.6 *Second price auction (revisited).* The second price auction was defined in Example 5.5, as an example of a VCG mechanism. For this mechanism, the allocation rule and payment rule are as follows:

$$x_r(\mathbf{b}) = \begin{cases} 1, & \text{if } b_r > b_s, \ s \neq r; \\ 0, & \text{if } b_r < \max_{s \neq r} b_s; \end{cases}$$

$$w_r(\mathbf{b}) = \begin{cases} \max_{s \neq r} b_s, & \text{if } x_r(\mathbf{b}) = 1; \\ 0, & \text{otherwise.} \end{cases}$$

(Note that we leave unspecified the tiebreaking rule when multiple agents have the same bids; any reasonable rule can be used.)

We have already shown that truthful declaration of one's valuation is a weak dominant strategy for this auction, so in particular, that also gives rise to Nash equilibria and Bayesian–Nash equilibria as well. Formally, for any configuration of valuations (v_1, \dots, v_R) , the bid vector where $b_r = v_r$ for all r is a Nash equilibrium. To study Bayesian–Nash equilibria, suppose that the valuations are drawn from some distribution F . Then the truthful declaration strategy vector \mathbf{a} where $a_r(v_r) = v_r$ for all r is a Bayesian–Nash equilibrium as well.

Example 5.7 *First price auction (revisited).* The first price auction was described in Example 3.7. For this mechanism, the allocation and payment rule are as follows:

$$x_r(\mathbf{b}) = \begin{cases} 1, & \text{if } b_r > b_s, \ s \neq r; \\ 0, & \text{if } b_r < \max_{s \neq r} b_s; \end{cases}$$

$$w_r(\mathbf{b}) = \begin{cases} b_r, & \text{if } x_r(\mathbf{b}) = 1; \\ 0, & \text{otherwise.} \end{cases}$$

As before, the definition in case of ties is left unspecified here, though any reasonable rule typically suffices.

It is straightforward to show that the first price auction has no (weak or strict) dominant strategies; we leave this as an exercise for the

reader. We thus turn our attention to Nash equilibria and Bayesian–Nash equilibria.

We start by considering Nash equilibria. Suppose that there are R bidders, with valuations $v_1 > v_2 > \dots > v_R$. In general, Nash equilibria may not exist, depending on the tiebreaking rule; informally, the highest valuation bidder should be able to win at a bid which is infinitesimally larger than the second highest valuation, but the allocation and payment rules have a discontinuity at that point. We resolve the issue by choosing a particular (unrealistic) tiebreaking rule: we suppose in case of ties that the item is awarded to the highest *valuation* bidder. (Of course such a rule cannot be implemented without knowing the valuations!) In this case it can be shown that the bid vector where $b_1 = v_2$, and $b_s = v_s$ for $s \geq 2$, is a Nash equilibrium: the highest valuation bidder wins the item, and pays the second highest valuation. (With a general tiebreaking rule, a similar insight can be obtained by turning to the ϵ -Nash equilibrium solution concept, where bidders are only asked to maximize payoff to within ϵ of their optimal utility.)

Next, we consider Bayesian–Nash equilibria. Suppose that agents' valuations are drawn from a uniform distribution on $[0, 1]$, independently across bidders. In this setting, we already showed in Example 3.7 that in the case of two bidders, a Bayesian–Nash equilibrium of the auction is where $a_r(v_r) = v_r/2$, i.e., each agent bids half her value. More generally, it can be shown that with R bidders, there exists a Bayesian–Nash equilibrium where each agent bids $a_r(v_r) = (R - 1)v_r/R$. Note that in these equilibria, players are not truthful, because when they win they have to pay what they bid. Thus agents are trading off the benefits of bidding higher (to increase the chance of winning) against the benefits of bidding lower (to lower the required payment upon winning).

Example 5.8 *All-pay auction.* We conclude with one other auction model; in an all-pay auction, the highest bidder wins, but *all* bidders must pay what they bid, regardless of whether they win or lose. All-pay auctions are commonly used to model scenarios where bidders' incur a *sunk cost* before they know the outcome of the auction. For example,

a simple model of competition for wireless channel might involve an all-pay auction, where the “bid” of a device is the power it uses to transmit. If we assume the highest power device “wins” the channel (i.e., successfully transmits), then we obtain an all-pay auction, since *all* devices incur a power cost (even those that are unable to transmit successfully).

Formally, the allocation and payment rules of an all-pay auction are defined as follows:

$$x_r(\mathbf{b}) = \begin{cases} 1, & \text{if } b_r > b_s, s \neq r; \\ 0, & \text{if } b_r < \max_{s \neq r} b_s; \end{cases}$$

$$w_r(\mathbf{b}) = b_r.$$

As before, we leave the tiebreaking rule unspecified in case of multiple bidders with the highest bid.

The all-pay auction is somewhat unusual: there are no (weak or strict) dominant strategies, and there do not exist any pure strategy Nash equilibria. There does exist a mixed strategy Nash equilibrium, however. We leave proofs of these facts as exercises for the reader.

The all-pay auction does possess Bayesian–Nash equilibria in general. In particular, suppose that agents’ valuations are drawn uniformly on $[0, 1]$, independently across bidders. Then it can be shown that with R bidders, it is a Bayesian–Nash equilibrium for every agent to bid according to $a_r(v_r) = (R - 1)v_r^R/R$. Notice that this bid is even lower than the equilibrium bid in the Bayesian–Nash equilibrium of the first price auction, since $v_r^R \leq v_r$; because bidders pay even when they lose, they become even more cautious. In particular (and in contrast to the first price auction), as $R \rightarrow \infty$, individual bidders’ bids converge to zero.

5.4.3 Discussion

Revenue equivalence and payoff equivalence. With independent private values, Bayesian–Nash equilibria of many auctions similar to those above exhibit two remarkable features known as *revenue equivalence* and *payoff equivalence*. In particular, consider any auction format where

(1) the highest bidder always wins; and (2) a bidder who bids zero earns zero utility (i.e., does not win and makes no payment). Also, assume that valuations are drawn i.i.d. from a distribution F . Finally, given R bidders, let \mathbf{a} be a Bayesian–Nash equilibrium of the auction that is *symmetric*, i.e., where $a_r(v_r) = s(v_r)$ for some function s that is continuous and strictly increasing with $s(0) = 0$; here symmetry means that all bidders use the same functional form to map their valuation to a bid.

In this setting, the *payoff equivalence theorem* states that the expected payoff to a bidder with valuation v in equilibrium is independent of the auction format, and in particular given by $\int_0^v F(v)^{R-1} dv$; and the *revenue equivalence theorem* states that the expected revenue to the auctioneer is independent of the auction format, and in particular the same as the expected revenue in the truthful equilibrium of the second price auction — i.e., the expected value of the second highest valuation. (We omit the proofs of these two results; see the endnotes for references and further reading.)

This pair of results are significant in the auction theory literature, because they demonstrate that key quantities — payoffs and revenue — are insensitive to auction format in a broad sense, provided that one focuses on symmetric Bayesian–Nash equilibria. In particular, from a practical standpoint, the two results are often used to suggest that the particulars of the format are inessential. Of course, in practice, there are many reasons to be skeptical of such an argument: the assumptions behind the two equivalence theorems are strong, since bidders must not only have i.i.d. private valuations, but they must also hold exactly the same *beliefs* about those valuations as well, and also must bid following the same functional strategy in equilibrium. Further, bidders must have only scalar valuations — such equivalence results typically break down with multidimensional valuations. Finally, it is critical that we focused only on a single shot auction; in dynamic settings where auctions are repeated over time, the equivalence theorems typically break down. Thus, rather than suggesting that auction format does not matter, these results focus our attention on those features of practical auction designs that *do* in fact have a significant impact on allocative efficiency and revenue.

Risk neutrality. When using Bayesian–Nash equilibria for auctions, one assumption that is central to our analysis is that players are *risk neutral*, i.e., they care only about their expected wealth. This assumption can be questioned, and indeed there has been significant work addressing models that relax this assumption. Typically, the assumption is relaxed by assuming that agents maximize expected utility, where their utility function is a function of final wealth. Unfortunately, the cost of this modification is often the loss of tractability; for example, the payoff and revenue equivalence theorems typically no longer hold.

Optimal auctions. The literature on optimal auctions typically concerns the design of auction formats that maximize an objective of direct interest to the auctioneer; the most common optimal auction design problems involve *revenue maximization*. In the Bayesian setting, the preceding results suggest that if the auctioneer wishes to maximize revenue, then something other than auction format should be varied.

Under some assumptions on the distribution of bidder valuations, it can be shown that revenue can be maximized if the auctioneer sets a *reserve price*. In an auction with a reserve price, the highest bidder wins the item only if her bid is greater than the reserve; otherwise, no one wins the item. It can be shown that if bidders' valuations are drawn i.i.d. from a distribution with monotone hazard rate (the uniform distribution is one such example), then the expected-revenue-maximizing auction is a second price auction with a positive reserve price.

Reserve prices seem puzzling at first glance, since in some situations the auctioneer is guaranteed to earn nothing; how, then, do reserves help increase expected revenue? The key insight is that reserve prices ensure that *high* valuation bidders sometimes pay much more than they would have paid without the reserve; in particular, if all other bidders' valuations are low, then the high valuation bidder receives quite a discount in a standard second price auction. With the reserve price, however, the auctioneer is able to extract much of this value for herself. This is an important insight from optimal auction design: the auctioneer wants to extract revenue from high valuation bidders; but since the auctioneer doesn't have *a priori* knowledge of the bidders' valuations, targeting the high valuation bidders typically means some revenue from low valuation bidders must be relinquished.

We conclude by noting that for the optimal auction problem to be well-formulated, an additional constraint is required: namely, that each agent will voluntarily choose to participate in the auction. This constraint, known as *individual rationality*, typically plays a critical role in mechanism design problems where the objective is not simply Pareto efficiency, but something else (e.g., revenue maximization). In the absence of such a constraint, the bidders can be “held hostage” and forced to pay arbitrarily high amounts — a clearly implausible market design. Indeed, informally, in many settings an optimal mechanism is found by maximizing the desired objective subject to the constraint that the resulting mechanism is both individually rational and incentive compatible; we see one example of this in our discussion of the principal-agent problem in the following section.

The role of the equilibrium concept. As already suggested by our analysis of equilibria in the three auction formats, there can be significant differences in our predictions for auction outcomes depending on which equilibrium concept we choose to work with. Here we emphasize that point by comparing revenues and efficiency in the first and second price auctions.

We start with dominant strategies. As noted above the second price auction has a weak dominant strategy (truthtelling), while the first price auction has no dominant strategy. If all agents tell the truth in the second price auction, then the efficient allocation is achieved: the highest valuation bidder receives the item. Further, the revenue to the auctioneer is the second highest valuation. Thus, a dominant strategy view of the two auction formats suggests the second price auction is preferable.

However, now consider Nash equilibria. It is straightforward to establish that there exist Nash equilibria of the second price auction where bidders *other* than the highest valuation bidder win the item. For example, suppose that with two bidders with valuations $v_1 > v_2$, bidder 1 bids $b_1 < v_2$, and bidder 2 bids $b_2 > v_1$. Then bidder 2 wins, and bidder 1 has no incentive to deviate. Admittedly, this equilibrium is somewhat strange, as it requires a potentially implausible threat from bidder 2 to bid *above* bidder 1’s valuation. However, it is worth noting that no such inefficient Nash equilibria exist for the first price

auction (with the tiebreaking rule introduced above): all Nash equilibria of the first price auction have the item allocated to the highest valuation bidder.

We now consider revenues at Nash equilibria. Even when agents are restricted to bid no more than their valuations, there exist Nash equilibria of the second price auction where the revenue to the auctioneer is zero: in particular, in the example above, it is a Nash equilibrium for bidder 1 to bid $b_1 = v_1$ and bidder 2 to bid $b_2 = 0$. On the other hand, in a first price auction (with the tiebreaking rule introduced above), the revenue to the auctioneer in a Nash equilibrium is never lower than the second highest valuation. (We leave a proof of this as an exercise for the reader.)

The preceding discussion suggests that if the equilibrium concept used is Nash equilibrium, then perhaps the first price auction appears somewhat more attractive than the second price auction. However, as we discussed in Chapter 3, justifying the use of Nash equilibria typically requires agents to have more information about each other than is required for justification of dominant strategies. More generally, this discussion illustrates the importance of understanding equilibrium concepts: the choice of equilibrium concept can have a first order impact on how an auction format performs. Therefore it is critical that the equilibrium concept be chosen to match the information structure in the application of interest as closely as possible.

Common values and the “winner’s curse. In all the examples we have considered so far, individuals have private values; in other words, their payoff depends only on a single quantity that is known to them. An alternative model is one with *common values*; in the most extreme case, a common value auction is one where the value of the item is the *same* to all bidders. Such auctions are most interesting when bidders do not know this common value *a priori*, but rather receive a noisy signal of what it might be.

In Bayesian–Nash equilibria of the auction formats discussed above with common values, it can be shown that players typically *underbid* relative to the bid they would make without the common value assumption (in a sense that can be made precise). This effect is known as the “winner’s curse”: by winning the auction in a common value setting,

a bidder instantly feels regret, because she knows that *all other bidders received signals that suggested a lower bid was optimal*. Therefore the winning bidder “wishes” she had bid lower. Since bidders are rational, they factor in this regret *ex ante* prior to submission of their bid, leading to a lower bid than they otherwise would make. Though not immediately apparent, the winner’s curse is another example of *adverse selection* at work (cf. Section 2.6): particularly in a setting with many bidders, the winning bidder knows that she was most likely poorly informed relative to the population of bidders at large.

Other auction formats. The auctions we have discussed thus far in this section have two distinguishing features. First, the space of bids available to an agent is the same as the space of valuations (the nonnegative real numbers in both cases); in other words, we could alternatively view these auctions as mechanisms where individuals are directly asked to reveal their valuation (as in VCG mechanisms). Such mechanisms — where the strategy space of an agent *is* the space of valuations — are known as *direct mechanisms*. A central result in optimal auction design, the *revelation principle*, argues that in finding an optimal mechanism it suffices to restrict attention to direct mechanisms, and for this reason much of optimal mechanism design focuses on direct mechanisms alone.

In practice, however, there may be restrictions that make direct mechanisms impractical; this is particularly the case in settings where valuation functions can be quite complex. For example, *combinatorial auctions* are market mechanisms where individuals bid over bundles of goods. With N goods, an individual bidder’s valuation function is a 2^N -dimensional object: one valuation for each possible bundle. Clearly working with exponentially large bids is prohibitive; as a consequence, these are settings where more reasonable indirect mechanisms are often favored.

Even in cases where direct mechanisms are practical, indirect mechanisms are often favored. For example, a well-known indirect mechanism for auctioning a single good as in our examples is an *ascending price auction* in which the auctioneer announces a price for the good and bidders respond if they are willing to buy. The auctioneer then increases the price of the good until there is only one bidder remaining.

This implements the same outcome as the second price auction, but is typically preferable in practice: it is “simpler” for bidders to understand, and as a consequence *discovery* of the correct clearing price is easier to ensure.

The second distinguishing feature of the auctions we consider is that they are *static*. In practice, most instances of auctions employed are actually *repeated* or *dynamic*. A prominent example is provided by the sponsored search auctions run by Google and Microsoft. These are markets where advertisers bid on keywords, and the market auctions off attention from visitors who search on those keywords (formally, the markets sell *clicks* from searching users). These are markets that are literally run on *every* search conducted — as a consequence, billions of auctions per day. Although a first-cut analysis of such dynamic markets typically focuses on understanding the static mechanism first, dynamic and repeated auctions can lead to significant challenges in understanding incentives. This remains a fruitful and active area of current research.

5.5 The Principal-Agent Problem

We conclude this chapter by discussing the canonical model at the foundation of the theory of contracts in economics, known as the *principal-agent problem*. In some sense, the principal-agent problem is at the heart of the study of incentives and mechanism design: in particular, how do we design mechanisms that incentivize others towards desired behavior? The principal-agent problem codifies this problem in a setting where one party (the “principal”) wishes to motivate another self-interested party (the “agent”) to take actions that are in the principal’s best interest. The principal’s vehicle is a *contract* that sets terms for how the agent will be paid; the principal-agent problem then concerns design of an optimal contract. Principal-agent problems abound throughout economics; some specific conceptual examples are, e.g., contracts between content providers (the principals) and network service providers (their agents); and even regulation imposed by the government (the principal) to encourage network service providers (the agents) to provide “neutral” service.

5.5.1 Methodology

A central feature of the principal-agent problem that we study is *moral hazard*: namely, in many settings, the principal does not have the ability to observe the *actions* of the agent, and thus cannot contract on the action itself. Rather, the principal only observes the eventual *outcome* (and her own utility as a result). As a consequence, the principal can only contract on the outcome. The main issue we investigate is whether the principal can design a contract only on the outcome that incentivizes the agent to take actions in the principal's best interest.

It should be clear that to do so, the principal needs to somehow align the agent's payoff with the outcome of the project itself. To make this intuition formal, we consider a simple model with the following features.

- (1) *Actions.* The action of the agent is a real number $a \geq 0$; we interpret a as "effort," with higher values meaning higher effort.
- (2) *Outcome.* We assume the outcome of the work is $y = a + \epsilon$, where ϵ is a random variable with zero mean that is initially unobservable to either the principal or the agent.
- (3) *Linear contract.* For simplicity, we restrict attention to settings where the payment from the principal to the agent is of the form $w = s + by$, where s and b are both parameters to be chosen.
- (4) *Payoffs.* We assume the principal wishes to maximize her expected payoff, and thus wants to maximize $E[y - w]$. The agent incurs a cost of effort $c(a)$ per unit effort expended; we assume $c(0) = 0$, and that c is strictly convex and strictly increasing over $a \geq 0$. The agent chooses an action to maximize her net expected payoff, which is $E[w - c(a)]$.
- (5) *Participation.* We allow the possibility that the agent may choose not to accept the contract; in this case the principal earns zero payoff, and the agent earns a *reservation utility* \bar{w} .

The sequence of events is as follows. The principal offers a linear contract to the agent (a particular choice of s and b), and the agent

chooses whether to accept or reject. If the agent accepts, she then chooses an action a . Finally, the noise ϵ — and hence the outcome y — are realized, yielding the payment to the agent and the payoff to the principal.

The principal's goal is to choose s and b to maximize her own payoff, assuming the agent will be self-interested. Formally, this can be seen as a maximization problem subject to two constraints: *individual rationality* (the agent will only participate if it is beneficial in expected payoff terms to do so); and *incentive compatibility* (conditional on participating, the agent will take an action to maximize her own expected payoff).

What contract should the principal choose? Observe that the utilities of both the principal and agent are quasilinear (cf. Section 2.4), and thus any (Pareto) efficient outcome must maximize the *sum* of their expected payoffs, i.e.:

$$\mathbb{E}[y - c(a)] = \mathbb{E}[a + \epsilon] - c(a) = a - c(a). \quad (5.9)$$

We refer to the preceding quantity as the *total expected surplus*. Note that by our assumptions on c , there exists a unique $a^* > 0$ that maximizes the preceding quantity.

To gain some intuition, we first consider the case where $\bar{w} = 0$, i.e., as long as the agent has a nonnegative expected payoff from signing the contract, she will participate. In this case, observe that the principal can offer the agent a contract with $b^* = 1$ and $s^* = c(a^*) - a^*$ (the negation of the maximum total expected surplus). If the agent now considers choosing action a^* , her expected wage will be $w^* = \mathbb{E}[s^* + b^*y] = c(a^*) - a^* + a^* = c(a^*)$, and she will earn expected payoff $s^* + b^*\mathbb{E}[y] - c(a^*) = 0$, and thus be exactly indifferent between participating and not participating. Further, in this contract, the principal earns an expected payoff of $\mathbb{E}[y - w^*] = a^* - c(a^*)$, the maximum possible total surplus. It is clear that the principal can do no better than this: the contracted wage needs to be at least enough to compensate the agent for the cost of action, and subject to this constraint, $a^* - c(a^*)$ is the maximum attainable expected payoff. A similar analysis can be carried out if $\bar{w} > 0$; in this case the optimal contract still has $b = 1$, but the principal must sacrifice some surplus to incentivize the agent to participate.

Observe that in this setting, despite being unable to observe the action of the agent directly, the principal is able to structure a contract that generates the best possible *ex ante* choice of action. This is sometimes called the “first-best” outcome (cf. Section 2.4). From a contractual standpoint, this is the best benchmark the principal can hope for.

5.5.2 Discussion

Moral hazard and adverse selection; signaling and screening. As noted above, principal agent problems typically exhibit *moral hazard*, i.e., the situation where there is hidden action as part of the model. This can be contrasted with *adverse selection*, where there is hidden information about an agent’s preferences. For example, the “market for lemons” discussed in Section 2.6 exhibits adverse selection: in that model, a buyer does not know in advance the quality of the car she purchases from a seller, and lower quality sellers are incentivized to take advantage of buyers’ lack of information. Taken together, moral hazard and adverse selection present two of the most significant classes of asymmetric information in economic models.

In settings with adverse selection, two kinds of principal-agent problems may arise. One is of the form described above, where a principal offers a contract structure before the agent makes a decision; in this case the goal is to structure the contract to reveal as much information about the agent as is necessary for the principal to maximize her utility. Such problems are called *screening* problems. By contrast, our discussion of the market for lemons should be viewed as a setting with *signaling*: for example, in eBay, the advertisement of a seller (the agent) is a signal of the quality of the good; based on this advertisement a buyer (the principal) must decide whether they want to bid for the good or not. The key distinguishing feature between screening and signaling is the order of moves: in a screening setting, the principal moves first; in a signaling setting, the agent moves first.

Losses of efficiency, contracting on the outcome, and risk aversion. In our model, both the principal and the agent are *risk neutral*: they

directly contract on the payoff. The model can change significantly if this assumption is violated; here we briefly discuss the consequences of a *risk averse* agent.

A risk averse agent can be modeled by assuming that the agent maximizes $E[U(w)] - c(a)$, where U is some strictly concave, strictly increasing function of the wage. In this case, note that the uncertainty in the wage (due to ϵ) is no longer eliminated. This induces a dilemma in incentivizing the agent: contracting directly on the outcome ($b = 1$) aligns incentives properly, but may force the agent to decline the contract as too risky; on the other hand, a contract that is flat rate ($b = 0$) removes risk from the agent's point of view, but suffers from poor control of incentives (since it is no longer tied to the outcome). The optimal contract in this case must do enough to assuage the agent's concerns about risk while also incentivizing her at the same time, and generally has $0 < b < 1$ (when optimizing among linear contracts). Further, and as a consequence of risk aversion of the agent, the first-best outcome is typically no longer sustainable. There is some loss of efficiency as the balance between managing incentives and managing risk enters into the principal's optimization problem.

5.6 Endnotes

Many of the topics in this chapter are covered in greater detail in standard textbooks on microeconomic theory; see, e.g., Mas-Colell et al. [24] or Varian [39]. Some specific references related to material in each section follow.

Pricing. The pricing of externalities was first formalized by Pigou [35]. The example of interference prices was drawn from the work of Huang et al. [16].

VCG mechanisms. For foundational reading, see the seminal papers of Vickrey [40], Clarke [8], and Groves [13]. See also Section 23 of the textbook by Mas-Colell et al. [24].

Auctions. A useful introduction to auction theory, including coverage of the payoff equivalence and revenue equivalence theorems, optimal auctions, and common values, is the text by Krishna [21].

For a discussion of Vickrey auctions in practice (and why they are rare), see Ausubel and Milgrom [3]. See also Johari and Tsitsiklis [17] for a discussion of the strategic complexity of VCG mechanisms.

The principal-agent problem. A comprehensive introduction to principal-agent problems and contract theory is provided by Bolton and Dewatripont [7].

Acknowledgments

We would like to gratefully acknowledge helpful suggestions received from Frank Kelly, Thanh Nguyen, Rakesh Vohra, and Jean Walrand. We would also like thank the editor-in-chief of this series, Tony Ephremides, and the publisher, James Finlay, for inviting us to prepare a manuscript on this topic, and for their encouragement (and patience!) over the time we took to complete the monograph. Finally, we acknowledge funding from the U.S. National Science Foundation and DARPA that supported our own research, as well as preparation of this manuscript.

References

- [1] G. Akerlof, “The market for lemons: Qualitative uncertainty and the market mechanism,” *The Quarterly Journal of Economics*, vol. 84, no. 3, pp. 488–500, 1970.
- [2] R. Aumann, “Agreeing to disagree,” *The Annals of Statistics*, vol. 4, no. 6, pp. 1236–1239, 1976.
- [3] L. M. Ausubel and P. Milgrom, “The lovely but lonely Vickrey auction,” in *Combinatorial Auctions*, (P. Cramton, Y. Shoham, and R. Steinberg, eds.), pp. 17–40, MIT Press, Chap. 3, 2006.
- [4] A. Banerjee, “A simple model of herd behavior,” *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, August 1992.
- [5] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, 2nd Edition, 1992.
- [6] S. Bikhchandani, D. Hirshleifer, and I. Welch, “A theory of fads, fashion, customs and cultural change as informational cascades,” *The Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, October 1992.
- [7] P. Bolton and M. Dewatripont, *Contract Theory*. MIT Press, 2004.
- [8] E. H. Clarke, “Multipart pricing of public goods,” *Public Choice*, vol. 11, pp. 17–33, 1971.
- [9] R. Coase, “The problem of social cost,” *Journal of Law and Economics*, vol. 3, pp. 1–44, October 1960.
- [10] D. Fudenberg and D. Levine, *The Theory of Learning in Games*. MIT Press, 1998.
- [11] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.
- [12] R. Gibbons, *Game Theory for Applied Economists*. Princeton University Press, 1992.

- [13] T. Groves, “Incentives in Teams,” *Econometrica*, vol. 41, no. 4, pp. 617–631, 1973.
- [14] J. Harsanyi, “Games with incomplete information played by Bayesian players, parts I–III,” *Management Science*, vol. 14, pp. 159–182, 320–334, 486–502, 1967.
- [15] J. Harsanyi, “Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points,” *International Journal of Game Theory*, vol. 2, no. 1, pp. 1–23, 1973.
- [16] J. Huang, R. Berry, and M. L. Honig, “Distributed interference compensation for wireless networks,” *IEEE on Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1074–1085, May 2006.
- [17] R. Johari and J. N. Tsitsiklis, “Efficiency of scalar-parameterized mechanisms,” *Operations Research*, vol. 57, no. 4, pp. 823–839, 2009.
- [18] F. P. Kelly, “Charging and rate control for elastic traffic,” *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.
- [19] F. P. Kelly, A. K. Maulloo, and D. K. Tan, “Rate control for communication networks: Shadow prices, proportional fairness, and stability,” *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [20] E. Koutsoupias and C. Papadimitriou, “Worst-Case equilibria,” in *Proceedings of 16th Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Springer, 1999.
- [21] V. Krishna, *Auction Theory*. Academic Press, 2009.
- [22] S. H. Low and D. E. Lapsley, “Optimization flow control, I: Basic algorithm and convergence,” *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–875, 1999.
- [23] G. Mailath and L. Samuelson, *Repeated Games and Reputations*. Oxford University Press, 2006.
- [24] A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic Theory*. Oxford University Press, 1995.
- [25] P. Milgrom and N. Stokey, “Information, trade and common knowledge,” *Journal of Economic Theory*, vol. 26, no. 1, pp. 17–27, 1982.
- [26] J. Mo and J. Walrand, “Fair end-to-end window-based congestion control,” *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, October 2000.
- [27] D. Monderer and L. Shapley, “Potential Games,” *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [28] R. Myerson, *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.
- [29] R. B. Myerson and M. Satterthwaite, “Efficient mechanisms for bilateral trading,” *Journal of Economic Theory*, vol. 29, pp. 265–281, 1983.
- [30] J. F. Nash, “Equilibrium points in N-person games,” *Proceedings of the National Academy of Sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [31] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [32] M. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, 1994.
- [33] G. Owen, *Game Theory*. Emerald Group, 1995.

- [34] C. Papadimitriou, “Algorithms, games, and the Internet,” in *Proceedings of ACM Symposium on the Theory of Computing*, pp. 749–753, 2001.
- [35] A. C. Pigou, *The Economics of Welfare*. London, United Kingdom: Macmillan and Co., 1932. (Originally Published in 1920).
- [36] T. Roughgarden and É. Tardos, “How bad is selfish routing?,” *Journal of the ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [37] S. Shakkottai and R. Srikant, “Network optimization and control,” *Foundations and Trends in Networking*, vol. 2, no. 3, pp. 271–379, 2007.
- [38] S. Shenker, “Fundamental design issues for the future Internet,” *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1176–1188, 1995.
- [39] H. R. Varian, *Microeconomic analysis*. New York: W.W. Norton and Co., 1992.
- [40] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.
- [41] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. John Wiley and Sons, 1944.
- [42] J. Walrand, “Economic models of communication networks,” in *Performance Modeling and Engineering*, (Z. Liu and C. Xia, eds.), pp. 57–89, Springer, Chap. 3, 2008.