

The Price of Anarchy and the Design of Scalable Resource Allocation Mechanisms

Ramesh Johari

Abstract

In this chapter, we study the allocation of a single infinitely divisible resource among multiple competing users. While we aim for efficient allocation of the resource, the task is complicated by the fact that users' utility functions are typically unknown to the resource manager. We study the design of resource allocation mechanisms that are approximately efficient (i.e., have a low price of anarchy), with low communication requirements (i.e., the strategy spaces of users are low dimensional).

Our main results concern the *proportional allocation mechanism*, for which a tight bound on the price of anarchy can be provided. We also show that in a wide range of market mechanisms that use a single market-clearing price, the proportional allocation mechanism minimizes the price of anarchy. Finally, we relax the assumption of a single market-clearing price, and show that by extending the class of Vickrey–Clarke–Groves mechanisms all Nash equilibria can be guaranteed to be fully efficient.

21.1 Introduction

This chapter deals with a canonical resource allocation problem. Suppose that a finite number of users compete to acquire a share of an infinitely divisible resource of fixed capacity. How should the resource be shared among the users? We will frame this problem as an economic problem: we assume that each user has a utility function that is increasing in the amount of the resource received, and then design a mechanism to maximize aggregate utility. In the absence of any strategic considerations, this is a simple optimization problem; however, if we assume that the agents are strategic, we need to design the resource allocation mechanisms to be robust to gaming behavior.

A central theme of this chapter is that *the price of anarchy* can be used as a design metric; i.e., “robust” allocation mechanisms are those that have a low price of anarchy. The present chapter is thus a bridge between two different themes of the book. The first theme is that of *optimal mechanism design* (Part II): given selfish agents, how do we successfully design mechanisms that nevertheless yield efficient outcomes? The

second theme is that of quantifying inefficiency (Part III): given a prediction of game theoretic behavior, how well does it perform relative to some efficient benchmark? In this chapter, we use the quantification of inefficiency as the “objective function” with which we will design optimal mechanisms. As we will see, for the resource allocation problems we consider, this approach yields surprising insights into the structure of optimal mechanisms.

The mechanisms we consider for resource allocation are motivated by constraints present in modern communication networks, and similar systems where communication is limited; this precludes use of the traditional Vickrey–Clarke–Groves mechanisms (Chapter 9), which require declaration of the entire utility function. If we interpret the single resource above as a communication link, then we view the mechanism as an allocation policy operating on that link. We wish to design mechanisms that, intuitively, impose low communication overhead on the overall system; throughout this chapter, that scalability constraint translates into the assumption that the players can use only *low-dimensional* (in fact, one-dimensional) strategy spaces.

The remainder of the chapter is organized as follows. In Section 21.2, we introduce the basic resource allocation model we will consider in this chapter, and then introduce a simple approach to allocating the fixed resource: the *proportional allocation mechanism*. In this mechanism, each user submits a bid, and receives a share of the resource in proportion to their bid. We analyze this model under both the assumption that users are price takers (i.e., that they do not anticipate the effect of their strategic decision on the price of the resource); and the assumption that users are price anticipators. The former case yields full efficiency, while in the latter we characterize the price of anarchy. In Section 21.3, we state and prove a theorem showing that in a nontrivial class of “scalable” market mechanisms (in the sense informally discussed above), the proportional allocation mechanism has the lowest price of anarchy (i.e., minimizes the efficiency loss) when users are price anticipating.

In all the mechanisms considered in the first two sections, players have one-dimensional strategy spaces, and the mechanism also only chooses a single price. Because of these constraints, even the highest performance mechanisms suffer a positive efficiency loss, as demonstrated in Section 21.3. In the final section of the chapter, we consider the implications of removing the “single price” constraint. We show in Section 21.4 that if we consider mechanisms with scalar strategy spaces, and allow the mechanism to choose one price per user of the resource, then in fact full efficiency is achievable at Nash equilibrium. The result involves extending the well-known class of Vickrey–Clarke–Groves (VCG) mechanisms to use only a scalar strategy space; for more on VCG mechanisms, see Chapter 9.

21.2 The Proportional Allocation Mechanism

Suppose that R users share a resource of capacity $C > 0$. Let d_r denote the amount allocated to user r . We assume that user r receives a *utility* equal to $U_r(d_r)$ if the allocated amount is d_r ; we assume that utility is measured in monetary units. We make the following assumptions on the utility function; we emphasize that *this assumption will be in force for the duration of the chapter*, unless otherwise mentioned.

Assumption 1 For each r , over the domain $d_r \geq 0$ the utility function $U_r(d_r)$ is concave, strictly increasing, and continuous; and over the domain $d_r > 0$, $U_r(d_r)$ 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 utility functions satisfying these conditions.

We note that we make rather strong differentiability assumptions here on the utility functions; these assumptions are primarily made to ease the presentation. It is possible to relax the differentiability assumptions (see Notes for details).

Given complete knowledge and centralized control of the system, a natural problem for the network manager to try to solve is the following optimization problem:

SYSTEM:

$$\text{maximize } \sum_r U_r(d_r) \tag{21.1}$$

$$\text{subject to } \sum_r d_r \leq C; \tag{21.2}$$

$$d_r \geq 0, \quad r = 1, \dots, R. \tag{21.3}$$

Note that the objective function of this problem is the utilitarian social welfare function (cf. Chapter 17); it becomes a reasonable objective if we assume that all utilities are measured in the same (monetary) units. Since the objective function is continuous and the feasible region is compact, an optimal solution $\mathbf{d} = (d_1, \dots, d_R)$ exists. If the functions U_r are strictly concave, then the optimal solution is unique, since the feasible region is convex.

In general, the utility functions are not available to the resource manager. As a result, we consider the following pricing scheme for resource allocation, which we refer to as the *proportional allocation mechanism*. Each user r gives a payment (also called a *bid*) of w_r to the resource manager; we assume $w_r \geq 0$. Given the vector $\mathbf{w} = (w_1, \dots, w_r)$, the resource manager chooses an allocation $\mathbf{d} = (d_1, \dots, d_r)$. We assume the manager treats all users alike—in other words, the network manager does not *price discriminate*. Each user is charged the same price $\mu > 0$, leading to $d_r = w_r/\mu$. We further assume that the manager always seeks to allocate the entire resource capacity C ; in this case, we expect the price μ to satisfy

$$\sum_r \frac{w_r}{\mu} = C.$$

The preceding equality can only be satisfied if $\sum_r w_r > 0$, in which case we have

$$\mu = \frac{\sum_r w_r}{C}. \tag{21.4}$$

In other words, if the manager chooses to allocate the entire resource, and does not price discriminate between users, then for every nonzero \mathbf{w} there is a *unique* price $\mu > 0$, which must be chosen by the network, given by the previous equation.

We can interpret this mechanism as a *market-clearing* process by which a price is set so that demand equals supply. To see this interpretation, note that when a user chooses a total payment w_r , it is as if the user has chosen a *demand function* $D(p, w_r) = w_r/p$

for $p > 0$. The demand function describes the quantity the user demands at any given price $p > 0$. The resource manager then chooses a price μ so that $\sum_r D(\mu, w_r) = C$, i.e., so that the aggregate demand equals the supply C . For the specific form of demand functions we consider here, this leads to the expression for μ given in (21.4). User r then receives an allocation given by $D(\mu, w_r)$, and makes a payment $\mu D(\mu, w_r) = w_r$. This interpretation will be further explored in Section 21.3, where we consider other market-clearing mechanisms for allocating a single resource in inelastic supply, with the users choosing demand functions from a family parameterized by a single scalar.

21.2.1 Price Taking Users and Competitive Equilibrium

In this section, we consider a *competitive equilibrium* between the users and the resource manager. A central assumption in the definition of competitive equilibrium is that each user does not anticipate the effect of their payment w_r on the price μ ; i.e., each user acts as a *price taker*. In this case, given a price $\mu > 0$, user r acts to maximize the following payoff function over $w_r \geq 0$:

$$P_r(w_r; \mu) = U_r\left(\frac{w_r}{\mu}\right) - w_r. \tag{21.5}$$

The first term represents the utility to user r of receiving a resource allocation equal to w_r/μ ; the second term is the payment w_r made to the manager. Observe that this definition is consistent with the notion that all utilities are measured in *monetary* units.

We now say a pair (\mathbf{w}, μ) with $\mathbf{w} \geq 0$ and $\mu > 0$ is a *competitive equilibrium* if users maximize their payoff as defined in (21.5), and the network “clears the market” by setting the price μ according to (21.4):

$$P_r(w_r; \mu) \geq P_r(\bar{w}_r; \mu) \quad \text{for } \bar{w}_r \geq 0, \quad r = 1, \dots, R; \tag{21.6}$$

$$\mu = \frac{\sum_r w_r}{C}. \tag{21.7}$$

The following theorem shows that under our assumptions, a competitive equilibrium always exists, and any competitive equilibrium maximizes aggregate utility.

Theorem 21.1 *There exists a competitive equilibrium (\mathbf{w}, μ) . In this case, the vector $\mathbf{d} = \mathbf{w}/\mu$ is an optimal solution to SYSTEM.*

PROOF The key idea in the proof is to use Lagrangian techniques to establish that optimality conditions for (21.6)–(21.7) are identical to the optimality conditions for the problem SYSTEM, under the identification $\mathbf{d} = \mathbf{w}/\mu$.

Observe that under Assumption 1, the payoff (21.5) is concave in w_r for any $\mu > 0$. Thus considering the first-order condition for maximization of $P_r(w_r; \mu)$ over $w_r \geq 0$, we conclude \mathbf{w} and μ are a competitive equilibrium if and only if

$$U'_r(d_r) = \mu, \quad \text{if } d_r > 0; \tag{21.8}$$

$$U'_r(0) \leq \mu, \quad \text{if } d_r = 0; \tag{21.9}$$

$$\sum_r d_r = C, \tag{21.10}$$

where $d_r = w_r/\mu$. A straightforward Lagrangian optimization shows that the preceding conditions are exactly the optimality conditions for the problem *SYSTEM*, so we conclude \mathbf{w} and μ are a competitive equilibrium if and only if $\mathbf{d} = \mathbf{w}/\mu$ is a solution to *SYSTEM* with Lagrange multiplier μ . Since at least one solution to *SYSTEM* must exist, the proof is complete. \square

Theorem 21.1 shows that under the assumption that the users of the resource behave as price takers, there exists a bid vector \mathbf{w} where all users have optimally chosen their bids w_r , with respect to the given price $\mu = \sum_r w_r/C$; and at this “equilibrium,” aggregate utility is maximized. However, when the price taking assumption is violated, the model changes into a game and the guarantee of Theorem 21.1 is no longer valid. We investigate this game in the following section.

21.2.2 Price Anticipating Users and Nash Equilibrium

We now consider an alternative model where the users of a single resource are *price anticipating*, rather than price takers. The key difference is that while the payoff function P_r takes the price μ as a fixed parameter in (21.5), price anticipating users will realize that μ is set according to (21.4), and adjust their payoff accordingly; this makes the model a game between the R players.

We use the notation \mathbf{w}_{-r} to denote the vector of all bids by users other than r ; i.e., $\mathbf{w}_{-r} = (w_1, w_2, \dots, w_{r-1}, w_{r+1}, \dots, w_R)$. Given \mathbf{w}_{-r} , each user r chooses w_r to maximize:

$$Q_r(w_r; \mathbf{w}_{-r}) = \begin{cases} U_r\left(\frac{w_r}{\sum_s w_s}C\right) - w_r, & \text{if } w_r > 0; \\ U_r(0), & \text{if } w_r = 0. \end{cases} \quad (21.11)$$

over nonnegative w_r . The second condition is required so that the resource allocation to user r is zero when $w_r = 0$, even if all other users choose \mathbf{w}_{-r} so that $\sum_{s \neq r} w_s = 0$. The payoff function Q_r is similar to the payoff function P_r , except that the user anticipates that the network will set the price μ according to (21.4). A *Nash equilibrium* of the game defined by (Q_1, \dots, Q_R) is a vector $\mathbf{w} \geq 0$ such that for all r :

$$Q_r(w_r; \mathbf{w}_{-r}) \geq Q_r(\bar{w}_r; \mathbf{w}_{-r}), \quad \text{for all } \bar{w}_r \geq 0. \quad (21.12)$$

Note that the payoff function in (21.11) may be discontinuous at $w_r = 0$, if $\sum_{s \neq r} w_s = 0$. This discontinuity may preclude existence of a Nash equilibrium; it is easy to see this in the case where the system consists of only a single user with a strictly increasing utility function. Nevertheless, as long as at least two users are competing, it is possible to show that a unique Nash equilibrium exists, by noting that such an equilibrium solves a version of the *SYSTEM* problem but with “modified” utility functions.

Theorem 21.2 *Suppose that $R > 1$. Then there exists a unique Nash equilibrium $\mathbf{w} \geq 0$ of the game defined by (Q_1, \dots, Q_R) , and it satisfies $\sum_r w_r > 0$. In*

this case, the vector \mathbf{d} defined by

$$d_r = \frac{w_r}{\sum_s w_s} C, \quad r = 1, \dots, R, \quad (21.13)$$

is the unique optimal solution to the following optimization problem:

GAME:

$$\text{maximize } \sum_r \hat{U}_r(d_r) \quad (21.14)$$

$$\text{subject to } \sum_r d_r \leq C; \quad (21.15)$$

$$d_r \geq 0, \quad r = 1, \dots, R, \quad (21.16)$$

where

$$\hat{U}_r(d_r) = \left(1 - \frac{d_r}{C}\right) U_r(d_r) + \left(\frac{d_r}{C}\right) \left(\frac{1}{d_r} \int_0^{d_r} U_r(z) dz\right). \quad (21.17)$$

PROOF The proof is similar to the proof of Theorem 21.1. The first key step is to note that at any Nash equilibrium, at least two components of \mathbf{w} must be positive; this follows from the payoff (21.11) (see Exercise 17.5). Given this fact, the payoff of each user w_r is strictly concave and continuous in w_r so that \mathbf{w} is a Nash equilibrium if and only if the following first-order conditions hold:

$$U'_r \left(\frac{w_r}{\sum_s w_s} C \right) \left(1 - \frac{w_r}{\sum_s w_s} \right) = \frac{\sum_s w_s}{C}, \quad \text{if } w_r > 0; \quad (21.18)$$

$$U'_r(0) \leq \frac{\sum_s w_s}{C}, \quad \text{if } w_r = 0. \quad (21.19)$$

Note that if we define $\rho = \sum_s w_s / C$ and $d_r = w_r / \rho$, then the preceding conditions can be rewritten as

$$\hat{U}'_r(d_r) = \rho, \quad \text{if } d_r > 0; \quad (21.20)$$

$$\hat{U}'_r(0) \leq \rho, \quad \text{if } d_r = 0; \quad (21.21)$$

$$\sum_r d_r = C. \quad (21.22)$$

Note that these are identical to (21.8)–(21.10), but for the modified objective function (21.14). Since the utility functions $\hat{U}_r(d_r)$ are strictly concave and continuous over $0 \leq d_r \leq C$, the preceding first-order conditions are sufficient optimality conditions for *GAME*. We conclude that \mathbf{w} is a Nash equilibrium if and only if $\sum_s w_s > 0$, and the resulting allocation \mathbf{d} solves the problem *GAME* with Lagrange multiplier $\rho = \sum_s w_s / C$. To conclude the proof, observe that *GAME* has a strictly concave and continuous objective function over a compact feasible region, and thus has a unique optimal solution. It is straightforward to verify that this implies uniqueness of the Nash equilibrium as well. \square

Note that the preceding theorem gives a form of “potential” for the game under consideration: the Nash equilibrium is characterized as the unique solution to a natural

optimization problem. However, the objective function for this optimization problem is not a true (exact or ordinal) potential for the game under consideration; this is because while the objective function (21.14) depends on *allocations*, the users' strategic decisions are *bids*. Notably, this observation is in sharp contrast to the potentials found for routing games in Chapter 18, or for network formation in Chapter 19. For example, we cannot use the objective function (21.14) to conclude that best response dynamics will converge for our game. Nevertheless, the optimization formulation will help us study the price of anarchy of the game in the following section. For later reference, we note the following corollary, which uses a *variational inequality* formulation of the preceding theorem.

Corollary 21.3 *Suppose that $R > 1$. Let \mathbf{w} be the unique Nash equilibrium of the game defined by (Q_1, \dots, Q_R) , and define \mathbf{d} according to (21.13). Then for any other vector $\bar{\mathbf{d}} \geq 0$ such that $\sum_r \bar{d}_r \leq C$, there holds:*

$$\sum_r \hat{U}'_r(d_r)(\bar{d}_r - d_r) \leq 0. \tag{21.23}$$

PROOF The stated condition follows easily from (21.20)–(21.22), the optimality conditions for the problem *GAME*. \square

21.2.3 Price of Anarchy

We let \mathbf{d}^S denote an optimal solution to *SYSTEM*, and let \mathbf{d}^G denote the unique optimal solution to *GAME*. We now investigate the price of anarchy of this system; i.e., how much utility is lost because the users are price anticipating? To answer this question, we must compare the utility $\sum_r U_r(d_r^G)$ obtained when the users fully evaluate the effect of their actions on the price, and the utility $\sum_r U_r(d_r^S)$ obtained by choosing the point that maximizes aggregate utility. (We know, of course, that $\sum_r U_r(d_r^G) \leq \sum_r U_r(d_r^S)$, by definition of \mathbf{d}^S .) As we show in the following theorem, the efficiency loss is exactly 25% in the worst case.

Theorem 21.4 *Suppose that $R > 1$. Suppose also that $U_r(0) \geq 0$ for all r . If \mathbf{d}^S is any optimal solution to *SYSTEM*, and \mathbf{d}^G is the unique optimal solution to *GAME*, then:*

$$\sum_r U_r(d_r^G) \geq \frac{3}{4} \sum_r U_r(d_r^S).$$

Furthermore, this bound is tight: for every $\epsilon > 0$, there exists a choice of R , and a choice of (linear) utility functions $U_r, r = 1, \dots, R$, such that

$$\sum_r U_r(d_r^G) \leq \left(\frac{3}{4} + \epsilon\right) \left(\sum_r U_r(d_r^S)\right).$$

PROOF Our proof will rely on the following constant β :¹

$$\beta = \inf_{U \in \mathcal{U}} \inf_{C > 0} \inf_{0 \leq d, \bar{d} \leq C} \frac{U(d) + \hat{U}'(d)(\bar{d} - d)}{U(\bar{d})}. \quad (21.24)$$

Recall the definition of \mathcal{U} in Assumption 1, and of \hat{U} in (21.17).

Our proof involves using Corollary 21.3 to prove that β is a tight bound on the efficiency of Nash equilibria. We first establish that $\beta \geq 3/4$. Note that in (21.24), the quotient is strictly larger than 1 if $d > \bar{d}$, and equal to 1 if $d = \bar{d}$. Thus in computing β we can assume that $d < \bar{d}$ in (21.24). We then have:

$$\begin{aligned} U(d) + \hat{U}'(d)(\bar{d} - d) &= U(d) + U'(d) \left(1 - \frac{d}{\bar{d}}\right) (\bar{d} - d) \\ &\geq U(d) + \left(1 - \frac{d}{\bar{d}}\right) (U(\bar{d}) - U(d)) \\ &\geq \left(\frac{d}{\bar{d}}\right)^2 U(\bar{d}) + \left(1 - \frac{d}{\bar{d}}\right) U(\bar{d}) \\ &\geq \frac{3}{4} U(\bar{d}). \end{aligned}$$

The first inequality follows since $\bar{d} \leq C$ and U is concave. The second inequality follows since U is concave and nonnegative and $d \leq \bar{d}$, so $U(d) \geq (d/\bar{d})U(\bar{d})$. Finally, the third inequality follows since $x^2 - x + 1$ is minimized at $x = 1/2$. It follows from (21.24) that $\beta \geq 3/4$.

Next, we show that for any $\delta > 0$, there exists an example where the ratio of Nash aggregate utility to maximum aggregate utility is at least $\beta + \delta$. Our approach is essentially the same as that in Example 17.6. Fix U , $d < \bar{d}$, and let $C = \bar{d}$. Consider the following example. Suppose that $R > 1$ users compete for the resource. Let user 1 have utility function $U_1 = U$, and suppose users $2, \dots, R$ have linear utility functions with slope $\hat{U}'(d)$; i.e., $U_r(d_r) = \hat{U}'(d)d_r = (U'(d)(1 - d/C))d_r$. Let \mathbf{d}^S denote an optimal solution to *SYSTEM* for this model; since one feasible solution involves allocating the entire resource \bar{d} to user 1, we must have $\sum_s U_s(d_s^S) \geq U(\bar{d})$. On the other hand, recall that at any Nash equilibrium at least two users have positive quantities; and since the Nash equilibrium is unique, we conclude that all users $2, \dots, R$ receive the same positive quantity. Thus as $R \rightarrow \infty$, we must have $d_r \downarrow 0$ for $r = 2, \dots, R$. From (21.20)–(21.21), it follows that the Nash price $\sum_s w_s/C$ must converge to $\hat{U}'(d)$ as $R \rightarrow \infty$. Thus, at the Nash equilibrium, user 1 receives an allocation $d + \epsilon$, and all other users receive an allocation $(1 - d - \epsilon)/(R - 1)$, where $\epsilon \rightarrow 0$ as $R \rightarrow \infty$. The total Nash utility thus converges to $U(d) + \hat{U}'(d)(\bar{d} - d)$. The limiting ratio of Nash aggregate utility to maximum aggregate utility is thus less than or equal to

$$\frac{U(d) + \hat{U}'(d)(\bar{d} - d)}{U(\bar{d})}.$$

¹ A slight subtlety arises in this definition if $U(\bar{x}) = 0$; however, in this latter case we can define β by only taking the infimum over $\bar{x} > 0$. This does not change any of the subsequent arguments.

We conclude that for any $\delta > 0$, there exists a game (Q_1, \dots, Q_R) in which the ratio of Nash aggregate utility to maximum aggregate utility is at most $\beta + \delta$. By considering the special case in which $U(\hat{d}) = \hat{d}$, $d = 1/2$, and $\bar{d} = 1$, the preceding construction yields a limiting efficiency ratio of exactly $3/4$. Combined with the previous argument that $\beta \geq 3/4$, it follows that in fact $\beta = 3/4$.

It remains to show that the bound holds for *every* resource allocation game. Here we simply apply the result of Corollary 21.3. Let (Q_1, \dots, Q_R) be a resource allocation game where users have utility functions (U_1, \dots, U_R) . Let \mathbf{d}^S be a solution to *SYSTEM*, and let \mathbf{d}^G be a solution to *GAME*. We have

$$\sum_s U_s(d_s^S) \leq \sum_s \frac{1}{\beta} (U_s(d_s^G) + \hat{U}'_s(d_s^G)(d_s^S - d_s^G)) \leq \frac{1}{\beta} \sum_s U_s(d_s^G).$$

The first inequality follows by the definition of β , and the second follows from Corollary 21.3. Since $\beta = 3/4$, this concludes the proof. \square

The preceding theorem shows that in the worst case, aggregate utility falls by no more than 25% when users are able to anticipate the effects of their actions on the price of the resource. Furthermore, this bound is essentially tight. In fact, it follows from the proof that the worst case consists of a resource of capacity 1, where user 1 has utility $U_1(d_1) = d_1$, and all other users have utility $U_r(d_r) \approx d_r/2$ (when R is large). As $R \rightarrow \infty$, at the Nash equilibrium of this game user 1 receives a quantity $d_1^G = 1/2$, while the remaining users uniformly split the quantity $1 - d_1^G = 1/2$ among themselves, yielding an aggregate utility of $3/4$. On the other hand, the maximum aggregate utility possible is clearly 1, achieved by allocating the entire resource to user 1.

21.3 A Characterization Theorem

In this chapter we ask an axiomatic question: Is the mechanism we have chosen “desirable” among a class of mechanisms satisfying certain “reasonable” properties? Defining desirability is the simpler of the two tasks: we consider a mechanism to be desirable if it minimizes efficiency loss when users are price anticipating. Importantly, we ask for this efficiency property *independent* of the characteristics of the market participants (i.e., their cost functions or utility functions). That is, the mechanisms we seek are those that perform well under broad assumptions on the nature of the preferences of market participants.

How do we define “reasonable” mechanisms? The most important condition we impose is that the strategy space of each market participant should be “simple,” which we interpret as *low dimensional*. Formally, we will focus on mechanisms for which the strategy space of each market participant is \mathbb{R}^+ ; i.e., each market participant chooses a scalar, which is a parameter that determines his demand function as input to the market-clearing mechanism. The primary motivation is that if we view such a mechanism to be useful for a communication network setting, information flow is limited; and in particular, we would like to implement a market with as little overhead as possible.

Thus keeping the strategy spaces of the users low dimensional is a reasonable goal.² We will show that under a specific set of mathematical assumptions, the proportional allocation mechanism in fact minimizes the worst-case efficiency loss when users are price anticipating.

The class of market mechanisms we will consider is defined as follows. A market mechanism must operate on a particular environment, defined by a triple (C, R, \mathbf{U}) : $C > 0$ denotes the capacity of the resource; $R > 1$ denotes the number of users sharing the resource; and $\mathbf{U} = (U_1, \dots, U_R)$ denotes the utility functions of the users, with $U_r \in \mathcal{U}$ (cf. Assumption 1). The following definition captures our notion of a market mechanism.

Definition 21.5 A *smooth market-clearing mechanism* is a differentiable function $D: (0, \infty) \times [0, \infty) \rightarrow \mathbb{R}^+$ such that for all $C > 0$, for all $R > 1$, and for all nonzero $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$, there exists a unique solution $p > 0$ to the following equation:

$$\sum_{r=1}^R D(p, \theta_r) = C.$$

We let $p_D(\boldsymbol{\theta})$ denote this solution.³

Note that the market-clearing price is undefined if $\boldsymbol{\theta} = \mathbf{0}$. As we will see below, when we formulate a game between users for a given mechanism D , we will assume that the payoff to all players is $-\infty$ if the composite strategy vector is $\boldsymbol{\theta} = \mathbf{0}$. Note that this is slightly different from the definition in Section 21.1, where the payoff is $U(0)$ to a player with utility function U who submits a strategy $\theta = 0$. We will discuss this distinction further later; we simply note for the moment that it does not affect the results of this section.

Our definition of a smooth market-clearing mechanism generalizes the demand function interpretation of the proportional allocation mechanism. Recall that for that mechanism, each user submits a demand function of the form $D(p, \theta) = \theta/p$, and the link manager chooses a price $p_D(\boldsymbol{\theta})$ to ensure that $\sum_{r=1}^R D(p, \theta_r) = C$. Thus, for this mechanism, we have $p_D(\boldsymbol{\theta}) = \sum_{r=1}^R \theta_r / C$ if $\boldsymbol{\theta} \neq \mathbf{0}$.

We now generalize *competitive equilibria* and *Nash equilibria* to this setting.

Definition 21.6 Given a utility system (C, R, \mathbf{U}) and a smooth market-clearing mechanism D , we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a *competitive equilibrium* if, for $\mu = p_D(\boldsymbol{\theta})$, there holds for all r :

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} [U_r(D(\mu, \bar{\theta}_r)) - \mu D(\mu, \bar{\theta}_r)]. \quad (21.25)$$

² Note that this notion is distinct from “single-parameter domains” as studied in Chapter 9; there it is the true valuations of the agents that are one-dimensional, whereas here the true valuations of the agents may be arbitrary functions. With one-dimensional strategy spaces, we restrict the ability of users to *communicate* information about their valuations to the mechanism.

³ Note that we suppress the dependence of this solution on C ; where necessary, we will emphasize this dependence.

Definition 21.7 Given a utility system (C, R, \mathbf{U}) and a smooth market-clearing mechanism D , we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a *Nash equilibrium* if there holds for all r :

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} Q_r(\bar{\theta}_r; \boldsymbol{\theta}_{-r}). \tag{21.26}$$

where

$$Q_r(\theta_r; \boldsymbol{\theta}_{-r}) = \begin{cases} U_r(D(p_D(\boldsymbol{\theta}), \theta_r)) - p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r), & \text{if } \boldsymbol{\theta} \neq \mathbf{0}; \\ -\infty, & \text{if } \boldsymbol{\theta} = \mathbf{0}. \end{cases} \tag{21.27}$$

Notice that the payoff Q_r is $-\infty$ if the composite strategy vector is $\boldsymbol{\theta} = \mathbf{0}$, since in this case no market-clearing price exists.

We are now ready to frame the specific class \mathcal{D} of market mechanisms we will consider in this section, defined as follows.

Definition 21.8 The class \mathcal{D} consists of all functions $D(p, \boldsymbol{\theta})$ such that the following conditions are satisfied:

- (i) D is a smooth market-clearing mechanism (cf. Definition 21.5).
- (ii) For all $C > 0$, and for all $U_r \in \mathcal{U}$, a user’s payoff is *concave if he is price anticipating*; i.e., for all R , and for all $\boldsymbol{\theta}_{-r} \in (\mathbb{R}^+)^R$, the function:

$$U_r(D(p_D(\boldsymbol{\theta}), \theta_r) - p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r))$$

is concave in $\theta_r > 0$ if $\boldsymbol{\theta}_{-r} = \mathbf{0}$, and concave in $\theta_r \geq 0$ if $\boldsymbol{\theta}_{-r} \neq \mathbf{0}$.

- (iii) For all $p > 0$, and for all $d \geq 0$, there exists a $\theta > 0$ such that $D(p, \theta) = d$.
- (iv) The demand functions are nonnegative; i.e., for all $p > 0$ and $\theta \geq 0$, $D(p, \theta) \geq 0$.

We pause here to briefly discuss the conditions in the previous definition. The second allows us to characterize Nash equilibria in terms of only first-order conditions. To justify this condition, we note that some assumption of quasiconcavity is generally used to guarantee existence of pure strategy Nash equilibria. The third condition ensures that given a price p and desired allocation $d \in [0, C]$, each player can make a choice of θ to guarantee precisely the allocation d . This is an “expressiveness” condition on the mechanism that ensures that all possible demands can be chosen at any market-clearing price. The last condition is a normalization condition, which ensures that regardless of the bid of a user, he is never required to *supply* some quantity of the resource (which would be the case if we allowed $D(p, \theta) < 0$). The following example gives a family of mechanisms that lie in \mathcal{D} .

Example 21.9 Suppose that $D(p, \theta) = \theta p^{-1/c}$, where $c \geq 1$. It is easy to check that this class of mechanisms satisfies $D \in \mathcal{D}$ for all choices of c ; when $c = 1$, we recover the proportional allocation mechanism of Section 21.2. The market-clearing condition yields that $p_D(\boldsymbol{\theta}) = (\sum_r \theta_r / C)^{1/c}$. Note that as a result, the *allocation* to user r at a nonzero vector $\boldsymbol{\theta}$ is

$$D(p_D(\boldsymbol{\theta}), \theta_r) = \frac{\theta_r}{\sum_s \theta_s} C.$$

In other words, regardless of the value of c , the market clearing allocations are chosen proportional to the bids. This remarkable fact is a special case of a more general result we establish below: all mechanisms in \mathcal{D} yield market-clearing allocations that are proportional to the bids; they differ only in the market-clearing price that is chosen. The exercises study the price of anarchy of the mechanisms defined in this example using an approach analogous to the proof of Theorem 21.4.

Our interest is in the worst-case ratio of aggregate utility at any Nash equilibrium to the optimal value of *SYSTEM*. Formally, for $D \in \mathcal{D}$ we define a constant $\rho(D)$ as follows:

$$\rho(D) = \inf \left\{ \frac{\sum_{r=1}^R U_r(D(p_D(\boldsymbol{\theta}), \theta_r))}{\sum_{r=1}^R U_r(d_r)} \mid C > 0, R > 1, \mathbf{U} \in \mathcal{U}^R, \right. \\ \left. \mathbf{d} \text{ solves } \textit{SYSTEM}, \text{ and } \boldsymbol{\theta} \text{ is a Nash equilibrium} \right\}.$$

Note that since all $U \in \mathcal{U}$ are strictly increasing and nonnegative, the aggregate utility $\sum_{r=1}^R U_r(d_r^S)$ is positive for any utility system (C, R, \mathbf{U}) with $C > 0$, and any optimal solution \mathbf{d}^S to *SYSTEM*. Note also that we are considering the ratio over *all* possible Nash equilibria, not just the best one for a given instance; thus, we are studying the price of anarchy, not the price of stability (cf. Chapter 17). However, Nash equilibria may not exist for some utility systems (C, R, \mathbf{U}) ; in this case we set $\rho(D) = -\infty$.

Our main result in this section is the following theorem.

Theorem 21.10 *Let $D \in \mathcal{D}$ be a smooth market-clearing mechanism. Then:*

- (i) *There exists a competitive equilibrium $\boldsymbol{\theta}$. Furthermore, for any such $\boldsymbol{\theta}$, the resulting allocation \mathbf{d} given by $d_r = D(p_D(\boldsymbol{\theta}), \theta_r)$ solves *SYSTEM*.*
- (ii) *There exists a concave, strictly increasing, differentiable, and invertible function $B : (0, \infty) \rightarrow (0, \infty)$ such that for all $p > 0$ and $\theta \geq 0$:*

$$D(p, \theta) = \frac{\theta}{B(p)}.$$

- (iii) *$\rho(D) \leq 3/4$, and this bound is met with equality if and only if $D(p, \theta) = \Delta\theta/p$ for some $\Delta > 0$.*

Before continuing to the proof of the theorem, we pause to make several critical comments about the result. Results (i) and (ii) of the theorem are a *characterization* of the types of mechanisms allowed by the constraints that define \mathcal{D} . In particular, notice that from (ii), for nonzero $\boldsymbol{\theta}$ we have

$$B(p_D(\boldsymbol{\theta})) = \frac{\sum_{r=1}^R \theta_r}{C}. \tag{21.28}$$

Thus we must have

$$D(p_D(\boldsymbol{\theta}), \theta_r) = \frac{\theta_r}{\sum_s \theta_s} C; \tag{21.29}$$

in other words, every mechanism in \mathcal{D} chooses allocations in proportion to the bids. As a result, we conclude that for a given vector θ , when the market clears, mechanisms in \mathcal{D} differ from the proportional allocation mechanism only in the market-clearing price—the allocation is the same. Result (iii) of the theorem is then a price of anarchy result that concerns mechanisms of this form.

We emphasize that the theorem here is distinguished from related work because the allocation rule (21.29) was not assumed in advance. Rather, the result here starts from a set of simple assumptions on the structure of mechanisms to be considered (the definition of the class \mathcal{D}), and uses them to *prove* that any mechanism in the class must lead to the allocation in (21.29). (See Notes for details.)

PROOF Throughout the proof we fix a particular mechanism $D \in \mathcal{D}$. Some computational details are left to the reader.

Step 1: A user's payoff is concave if he is price taking. In other words, we will show that for all $U \in \mathcal{U}$ and for all $p > 0$, $U(D(p, \theta)) - pD(p, \theta)$ is concave in θ . The key idea is to use a limiting regime where capacity grows large, so that users that are price anticipating effectively become price taking.

Formally, we first observe that since D must possess a unique market-clearing price regardless of the value of C , $D(p, \theta)$ must be strictly monotonic in p (for fixed $\theta > 0$) where it is nonzero, and either (1) $D(p, \theta)$ is nondecreasing in p for all $\theta > 0$, or (2) $D(p, \theta)$ is nonincreasing in p for all $\theta > 0$.

To complete the proof of this step, fix $\mu > 0$, and fix $\theta > 0$. Now consider a limit where $R \rightarrow \infty$, and $C^R = RD(\mu, \theta)$ is the capacity in the R 'th system. It is straightforward to check that if the $R - 1$ users $2, \dots, R$ submit strategy θ , and the first user submits strategy θ' , then the resulting market-clearing price p_D converges to μ as $R \rightarrow \infty$, regardless of the value of θ' . This step uses the fact that either (1) or (2) above holds. Applying the fact that player 1's payoff must be concave when he is price anticipating and taking limits as $R \rightarrow \infty$, it follows that player 1's payoff is concave when he is price taking for any fixed price $\mu > 0$.

Step 2: There exists a positive function B such that $D(p, \theta) = \theta/B(p)$ for $p > 0$ and $\theta \geq 0$. By Step 1, a player's payoff is concave when he is price taking. By appropriately choosing a linear utility function with very large slope and very small slope, it follows that $D(p, \theta)$ must be concave and convex, respectively, in θ for a given $p > 0$. Thus for fixed $p > 0$, $D(p, \theta)$ is an affine function of θ . Conditions 3 and 4 in Definition 21.8 then imply that the constant term must be zero, while the coefficient of the linear term is positive; thus, $D(p, \theta) = \theta/B(p)$ for some positive function $B(p)$.

Before continuing, we note that the previous step already implies the remarkable fact that for any mechanism $D \in \mathcal{D}$, the allocation at the market-clearing price is made in proportion to the bids θ . This follows from the discussion following (21.28) above.

Step 3: For all utility systems (C, R, \mathbf{U}) , there exists a competitive equilibrium, and it is fully efficient. This step follows primarily because of Condition 3 in Definition 21.8: given a price μ , a user can first determine his optimal choice of quantity, and then choose a parameter θ to express this choice. Formally, suppose that $\mu = p_D(\theta)$, and (21.25) holds. Let $d_r = D(\mu, \theta_r)$; then (21.25) implies that the necessary conditions (21.8)–(21.9) hold; these are also sufficient because of Step 1. Furthermore, market clearing implies (21.10) holds. Thus any competitive equilibrium is fully efficient. Existence follows by letting \mathbf{d}^S be a solution to *SYSTEM* with Lagrange multiplier μ , and choosing $\theta_r = d_r / B(\mu)$.

Step 4: For all $R > 1$ and $\theta_{-r} \in (\mathbb{R}^+)^{R-1}$, the functions $D(p_D(\theta), \theta_r)$ and $-p_D(\theta)D(p_D(\theta), \theta_r)$ are concave in $\theta_r > 0$ if $\theta_{-r} = \mathbf{0}$, and concave in $\theta_r \geq 0$ if $\theta_{-r} \neq \mathbf{0}$. As in Step 2, this conclusion follows by considering linear utility functions with very large and very small slope, respectively.

Step 5: B is an invertible, differentiable, strictly increasing, and concave function on $(0, \infty)$. We immediately see that B must be invertible on $(0, \infty)$; it is clearly onto, as the right-hand side of (21.28) can take any value in $(0, \infty)$. Furthermore, uniqueness of the market-clearing price in (21.28) requires that B is one-to-one as well, and hence invertible. Since D is differentiable, B must be differentiable as well. Let Φ denote the differentiable inverse of B on $(0, \infty)$; we will show Φ is strictly increasing and convex.

Let

$$w_r(\theta) = p_D(\theta)D(p_D(\theta), \theta_r) = \Phi \left(\frac{\sum_{s=1}^R \theta_s}{C} \right) \left(\frac{\theta_r}{\sum_{s=1}^R \theta_s} C \right). \quad (21.30)$$

By Step 4, $w_r(\theta)$ is convex in $\theta_r > 0$. By considering strategy vectors θ for which $\theta_{-r} = \mathbf{0}$, it follows that Φ is convex. Finally, the fact that Φ is strictly increasing follows by differentiating twice and considering the limit where $\theta_r \rightarrow 0$, while keeping θ_{-r} constant and nonzero.⁴ This establishes the desired facts regarding B .

Step 6: Let (C, R, \mathbf{U}) be a utility system. A vector $\theta \geq \mathbf{0}$ is a Nash equilibrium if and only if at least two components of θ are nonzero, and there exists a nonzero vector $\mathbf{d} \geq \mathbf{0}$ and a scalar $\mu > 0$ such that $\theta_r = \mu d_r$ for all r , $\sum_{r=1}^R d_r = C$, and the following conditions hold:

$$U'_r(d_r) \left(1 - \frac{d_r}{C} \right) = \Phi(\mu) \left(1 - \frac{d_r}{C} \right) + \mu \Phi'(\mu) \left(\frac{d_r}{C} \right), \quad \text{if } d_r > 0; \quad (21.31)$$

$$U'_r(0) \leq \Phi(\mu), \quad \text{if } d_r = 0. \quad (21.32)$$

In this case $d_r = D(p_D(\theta), \theta_r)$, $\mu = \sum_{r=1}^R \theta_r / C$, and $\Phi(\mu) = p_D(\theta)$. Further, there exists a unique Nash equilibrium. The proof of this step is similar to the

⁴ While the most direct argument uses twice differentiability of Φ , it is possible to make a similar argument even if Φ is only once differentiable, by arguing only in terms of increments of Φ .

proof of Nash equilibrium characterization in Theorem 21.2; we omit the details, and refer the reader to the Notes section.

Step 7: For any $\epsilon > 0$, there exists a utility systems (C, R, \mathbf{U}) such that at any Nash equilibrium θ , the aggregate utility is no more than $3/4 + \epsilon$ of the maximal aggregate utility. Consider a utility system with the following properties. Let $C = 1$. Fix $\mu > 0$, and let $U_1(d_1) = Ad_1$, where $A > \Phi(\mu)$. We will search for a solution to the Nash conditions (21.31) to (21.32) with market-clearing price $\Phi(\mu)$.

We start by calculating d_1 by assuming it is nonzero, and applying (21.31):

$$d_1 = \frac{(A - \Phi(\mu))C}{A - \Phi(\mu) + \mu\Phi'(\mu)}. \tag{21.33}$$

In the spirit of the proof of Theorem 21.4, we will now choose users $2, \dots, R$ to have identical linear utility functions, with slopes less than A . As we will see, this will be possible if R is large enough.

Formally, let $d = (C - d_1)/(R - 1)$, and (cf. (21.31)) define

$$\alpha = \frac{\Phi(\mu)C + (\mu\Phi'(\mu) - \Phi(\mu))d}{C - d}. \tag{21.34}$$

Let $U_r(d_r) = \alpha d_r$ for $r = 2, \dots, R$. Note that if

$$\frac{C}{R} \leq \frac{(A - \Phi(\mu))C}{A - \Phi(\mu) + \mu\Phi'(\mu)}, \tag{21.35}$$

then $\alpha \leq A$. This guarantees d_1 must be nonzero at any Nash equilibrium, so that the computation in (21.33) is valid. In turn, letting $d_r = d$ for $r = 2, \dots, R$, this implies that (d_1, \dots, d_R) and μ are a valid solution to (21.31)–(21.32), when users have utility functions U_1, \dots, U_R .

Now consider the limiting ratio of Nash aggregate utility to maximal aggregate utility, as $R \rightarrow \infty$. We have $d \rightarrow 0$, so $\alpha \rightarrow \Phi(\mu)$. Furthermore, regardless of R a solution to *SYSTEM* is to allocate the entire resource to user 1, so the maximal aggregate utility is AC . Thus the limiting ratio of Nash aggregate utility to maximal aggregate utility becomes

$$\frac{(A - \Phi(\mu))}{A - \Phi(\mu) + \mu\Phi'(\mu)} + \left(1 - \frac{(A - \Phi(\mu))}{A - \Phi(\mu) + \mu\Phi'(\mu)}\right) \left(\frac{\Phi(\mu)}{A}\right). \tag{21.36}$$

We now want to find the choices of A and μ which minimize this value.

For notational simplicity, we define $x = \Phi(\mu)/A$, and $\Psi(\mu) = \mu\Phi'(\mu)/\Phi(\mu)$. Note that given the convexity and invertibility of Φ , we have $\Psi(\mu) \geq 1$. Then (21.36) is equivalent to

$$F(x; \mu) = \frac{(1 - x)^2}{1 + (\Psi(\mu) - 1)x} + x. \tag{21.37}$$

It is straightforward to establish that the preceding expression is strictly convex in x for fixed μ . Let $G(\Psi(\mu))$ denote the minimal value of $F(x; \mu)$ for $x \in (0, 1)$;




Figure 21.1. The function $G(\Psi)$ defined in (21.38). Note that $G(\Psi)$ is strictly decreasing, with $G(1) = 3/4$.

by differentiating, it follows that $G(\Psi)$ is defined for $\Psi \geq 1$ according to

$$G(\Psi) = \begin{cases} \frac{3}{4}, & \text{if } \Psi = 1; \\ \frac{2\Psi^2 - 3\Psi\sqrt{\Psi} + \sqrt{\Psi}}{(\Psi - 1)^2\sqrt{\Psi}}, & \text{if } \Psi > 1. \end{cases} \quad (21.38)$$

The function G is plotted in Figure 21.1. It is straightforward to verify that $G(\Psi)$ is continuous and strictly decreasing for $\Psi \geq 1$ so that the worst-case example is given by finding $\mu > 0$ such that $\Psi(\mu)$ is maximized. Furthermore, it is straightforward to check that $G(\Psi) \leq 3/4$, establishing the required claim.

Step 8: For any mechanism other than the proportional allocation mechanism, the worst-case efficiency is strictly lower than 3/4. For the proportional allocation mechanism, we have $\Psi(\mu) = 1$, and we have already established that the efficiency ρ is exactly 3/4. On the other hand, it is straightforward to check that if $B(p)$ is nonlinear, then the maximal value of $\Psi(\mu)$ in the preceding step is strictly greater than 1; and in this case $G(\Psi(\mu))$ is strictly less than 3/4. Thus there exists a game with efficiency ratio strictly lower than 3/4 for such a mechanism. This completes the proof. \square

We make several comments regarding the proof. First, notice that every mechanism in the described class allocates in *proportion* to the bids of the players; in this sense all mechanisms in \mathcal{D} are “proportional allocation mechanisms.” However, the efficiency loss is minimized exactly when this mechanism charges each user exactly their bid. Second, it is possible to show that the bound constructed in Steps 7–8 of the proof is

in fact a tight bound on the price of anarchy of the mechanisms under consideration; it is possible to reformulate this bound so that it depends only on the *elasticity* of the function $B(p)$, i.e., the quantity $\inf_{p>0} pB'(p)/B(p)$. (This is not surprising, since $\Psi(\mu)$ is the elasticity of the function Φ , which is the inverse of B .) It is surprising that the price of anarchy of a general class of such mechanisms can be reduced to this parsimonious calculation.

Finally, we note one potentially undesirable feature of the family of market-clearing mechanisms considered: the payoff to user r is defined as $-\infty$ when the composite strategy vector is $\theta = 0$ (cf. (21.27)). This definition is required because when the composite strategy vector is $\theta = 0$, a market-clearing price may not exist. One possible remedy is to restrict attention instead to mechanisms where $D(p, \theta) = 0$ if $\theta = 0$, for all $p \geq 0$; in this case we can *define* $p_D(\theta) = 0$ if $\theta = 0$, and let the payoff to user r be $U_r(0)$ if $\theta_r = 0$. This condition amounts to a “normalization” on the market-clearing mechanism. It is possible to show that this modification does not alter the conclusion of Theorem 21.10.

21.4 The Vickrey–Clarke–Groves Approach

The mechanisms we considered in the last section had several restrictions placed on them; chief among these are that (1) users are restricted to using “simple” strategy spaces and (2) the mechanism uses only a single price to clear the market. On the other hand, one could consider both generalizations where users are allowed to use more complex strategies, perhaps declaring their entire utility function to the market; and also, where price discrimination is allowed so that each user is charged a personalized per-unit price for the resource.

The best known solution employing both these generalizations is the VCG approach to eliciting utility information (see Notes, and Chapter 9). Such mechanisms allow users to declare their entire utility functions, and then charge users individualized prices so that they have the incentive to truthfully declare their utilities. We review VCG mechanisms in Section 21.4.1.

In this section we are interested in deciding whether the same outcome can be realized preserving restriction (1) above, but removing restriction (2): that is, can mechanisms with “simple” strategy spaces that employ price discrimination achieve full efficiency? In Section 21.4.2 we present an alternate class of mechanisms, inspired by the VCG class, in which users only submit scalar strategies to the mechanism; we call such mechanisms *scalar strategy VCG* (SSVCG) mechanisms. We show that these mechanisms have desirable efficiency properties. In particular, we establish existence of an efficient Nash equilibrium, and under an additional condition, we also establish that all Nash equilibria are efficient.

21.4.1 VCG Mechanisms

In the VCG class of mechanisms, the basic approach is to let the strategy space of each user r be the set \mathcal{U} of possible utility functions, as defined in Assumption 1, and structure the payments made by each user so that the payoff of each user r has the same

form as the objective function in *SYSTEM*, (21.1). As VCG mechanisms have been introduced in Chapter 9, we only use this section to fix notation for our subsequent discussion. For each r , we use \tilde{U}_r to denote the declared utility function of user r , and use $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_R)$ to denote the vector of declared utilities.

Suppose that user r receives an allocation d_r , but has to make a payment t_r ; we use the notation t_r to distinguish from the bid w_r of Section 21.2. Then the payoff to user r is

$$U_r(d_r) - t_r.$$

On the other hand, the social objective (21.1) can be written as

$$U_r(d_r) + \sum_{s \neq r} U_s(d_s).$$

Given a vector of declared utility functions $\tilde{\mathbf{U}}$, a VCG mechanism chooses the allocation $\mathbf{d}(\tilde{\mathbf{U}})$ as an optimal solution to *SYSTEM* for the declared utility functions $\tilde{\mathbf{U}}$. For simplicity, let $\mathcal{X} = \{\mathbf{d} \geq 0 : \sum_r d_r \leq C\}$; this is the feasible region for *SYSTEM*. Then for a VCG mechanism, we have

$$\mathbf{d}(\tilde{\mathbf{U}}) \in \arg \max_{\mathbf{d} \in \mathcal{X}} \sum_r \tilde{U}_r(d_r). \quad (21.39)$$

The payments are structured so that

$$t_r(\tilde{\mathbf{U}}) = - \sum_{s \neq r} \tilde{U}_s(d_s(\tilde{\mathbf{U}})) + h_r(\tilde{\mathbf{U}}_{-r}). \quad (21.40)$$

Here h_r is an arbitrary function of the declared utilities of users other than r . In general, we note that mechanisms of this form do not use a single price to clear the market; i.e., the per-unit price paid by user r , $t_r(\tilde{\mathbf{U}})/d_r(\tilde{\mathbf{U}})$, will not be the same for all users. (See also Exercise 21.3.)

For our purposes, the interesting feature of the VCG mechanism is that there exists a dominant strategy equilibrium that elicits the true utility functions from the users, and in turn (because of the definition of $\mathbf{d}(\tilde{\mathbf{U}})$) chooses an efficient allocation. (See Chapter 9 for a formal statement of these results, where it is shown that the VCG mechanism is *incentive compatible*.) In the next section, we explore a class of mechanisms inspired by the VCG mechanisms, but with limited communication requirements.

21.4.2 Scalar Strategy VCG Mechanisms

We now consider a class of mechanisms where each user's strategy is a submitted utility function (as in the VCG mechanisms) except that users are allowed only to choose from a given single parameter family of utility functions. One cannot expect such mechanisms to have efficient dominant strategy equilibria, and we will focus instead on the efficiency properties of the resulting Nash equilibria.

Formally, *scalar strategy VCG* (SSVCG) mechanisms allow users to choose from a given family of utility functions $\overline{U}(\cdot; \theta)$, parameterized by $\theta \in (0, \infty)$.⁵ We make the following assumptions about this family.

Assumption 2:

- (i) For every $\theta > 0$, the function $\overline{U}(\cdot; \theta) : d \mapsto \overline{U}(d; \theta)$ belongs to \mathcal{U} (i.e., it is concave, strictly increasing, continuous, and differentiable), and is also strictly concave.
- (ii) For every $\gamma \in (0, \infty)$ and $d \geq 0$, there exists a $\theta > 0$ such that $\overline{U}'(d; \theta) = \gamma$.⁶

Given θ , the mechanism chooses $\mathbf{d}(\theta)$ such that

$$\mathbf{d}(\theta) = \arg \max_{\mathbf{d} \in \mathcal{X}} \sum_r \overline{U}(d_r; \theta_r). \tag{21.41}$$

Since $\overline{U}(\cdot; \theta_r)$ is strictly concave for each r , the solution $\mathbf{d}(\theta)$ is uniquely defined. (Note the similarity between (21.39) and (21.41).)

By analogy with the expression (21.40), the monetary payment by user r is

$$t_r(\theta) = - \sum_{s \neq r} \overline{U}(d_s(\theta); \theta_s) + h_r(\theta_{-r}). \tag{21.42}$$

Here h_r is a function that depends only on the strategies $\theta_{-r} = (\theta_s, s \neq r)$ submitted by the users other than r . While we do not advocate any particular choice of h_r , a natural candidate is to define $h_r(\theta_{-r}) = \sum_{s \neq r} \overline{U}(d_s(\theta_{-r}); \theta_s)$, where $vd(\theta_{-r})$ is the aggregate utility maximizing allocation excluding user r . This leads to a natural scalar strategy analogue of the Clarke pivot mechanism (cf. Chapter 9).

Given h_r , the payoff to user r is

$$P_r(d_r(\theta), t_r(\theta)) = U_r(d_r(\theta)) + \sum_{s \neq r} \overline{U}(d_s(\theta); \theta_s) - h_r(\theta_{-r}).$$

A strategy vector θ is a *Nash equilibrium* if no user can profitably deviate through a unilateral deviation, i.e., if for all users r there holds:

$$P_r(d_r(\theta), t_r(\theta)) \geq P_r(d_r(\theta'_r, \theta_{-r}), t_r(\theta'_r, \theta_{-r})), \text{ for all } \theta'_r > 0. \tag{21.43}$$

We start with the following key lemma, proven using an argument analogous to the proof that truth-telling is a dominant strategy equilibrium of the VCG mechanism (see Chapter 9).

⁵ Note that, by contrast with Section 21.3, the choice of bid θ by a user indexes a utility function, rather than a demand function. However, this is not particularly crucial: if a user with utility function U maximizes $U(d) - pd$ (i.e., the user acts as a price taker), the solution yields the demand function $D(p) = (U')^{-1}(p)$. Up to additive constant, the utility function and demand function can be recovered from each other. Thus, equivalently, we could define SSVCG mechanisms where users submit demand functions from a parameterized class. We define our SSVCG mechanisms according to Assumption 2 to maintain consistency with the definition of VCG mechanisms in Section 21.4.1, as well as in Chapter 9.

⁶ Since we do not assume differentiability with respect to θ , the only differentiation of \overline{U} is with respect to the first coordinate d , and $\overline{U}'(d; \theta)$ will always stand for the derivative with respect to d .

Lemma 21.11 *Then the vector θ is a Nash equilibrium of the SSVCG mechanism if and only if for all r :*

$$\mathbf{d}(\theta) \in \arg \max_{\mathbf{d} \in \mathcal{X}} \left[U_r(d_r) + \sum_{s \neq r} \bar{U}(d_s; \theta_s) \right]. \quad (21.44)$$

PROOF Fix a user r . Since θ_r does not affect h_r , from (21.43) user r will choose θ_r to maximize the following effective payoff:

$$U_r(d_r(\theta)) + \sum_{s \neq r} \bar{U}(d_s(\theta); \theta_s). \quad (21.45)$$

The optimal value of the objective function in (21.44) is certainly an upper bound to user r 's effective payoff (21.45). Thus, given a vector θ , if (21.44) is satisfied for all users r , then (21.43) holds for all users r , and we conclude θ is a Nash equilibrium.

Conversely, given a vector θ , suppose that (21.44) is not satisfied for some user r . We will show θ cannot be a Nash equilibrium. Since \mathcal{X} is compact, an optimal solution exists to the problem in (21.44) for user r ; call this optimal solution \mathbf{d}^* . The vector \mathbf{d}^* must satisfy the first-order optimality conditions (21.8)–(21.10), which only involve the first derivatives $U'_r(d_r^*)$ and $\bar{U}'(d_s^*; \theta_s)$, $s \neq r$. Suppose now that user r chooses $\theta'_r > 0$ such that $\bar{U}'(d_r^*; \theta'_r) = U'_r(d_r^*)$. Then, \mathbf{d}^* also satisfies the optimality conditions for the problem (21.41). Since $\mathbf{d}(\theta'_r, \theta_{-r})$ is the unique optimal solution to (21.41) when the strategy vector is (θ'_r, θ_{-r}) , we must have $\mathbf{d}(\theta'_r, \theta_{-r}) = \mathbf{d}^*$. Thus we have

$$\begin{aligned} P_r(d_r(\theta), t_r(\theta)) &< U_r(d_r^*) + \sum_{s \neq r} \bar{U}(d_s^*; \theta_s) + h_r(\theta_{-r}) \\ &= U_r(d_r(\theta'_r, \theta_{-r})) + \sum_{s \neq r} \bar{U}(d_s(\theta'_r, \theta_{-r}); \theta_s) + h_r(\theta_{-r}) \\ &= P_r(d_r(\theta'_r, \theta_{-r}), t_r(\theta'_r, \theta_{-r})). \end{aligned}$$

(The first inequality follows by the assumption that (21.44) is not satisfied for user r .) We conclude that (21.43) is violated for user r , so θ is not a Nash equilibrium. \square

The following corollary states that there exists a Nash equilibrium which is efficient. Furthermore, at this efficient Nash equilibrium, all users truthfully reveal their utilities in a local sense: each user r chooses θ_r so that the declared marginal utility $\bar{U}'(d_r(\theta); \theta_r)$ is equal to the true marginal utility $U'_r(d_r(\theta))$.

Corollary 21.12 *For any SSVCG mechanism, there exists an efficient Nash equilibrium θ defined as follows: Let \mathbf{d}^S be an optimal solution to SYSTEM. Each user r chooses θ_r so that $\bar{U}'(d_r^S; \theta_r) = U'_r(d_r^S)$. The resulting allocation satisfies $\mathbf{d}(\theta) = \mathbf{d}^S$.*

PROOF By Assumption 2, each user r can choose θ_r so that $\bar{U}'(d_r^S; \theta_r) = U'_r(d_r^S)$. For this vector θ , it is clear that $\mathbf{d}(\theta) = \mathbf{d}^S$, since the optimal solution to (21.41) is uniquely determined, and the optimality conditions for (21.41) involve only the first derivatives $\bar{U}'(d_r(\theta); \theta_r)$. By the same argument it also follows that \mathbf{d}^S is an optimal solution in (21.44). Since $\mathbf{d}(\theta) = \mathbf{d}^S$, we conclude that (21.44) is satisfied for all r , and thus θ is a Nash equilibrium. \square

We note that, as in classical VCG mechanisms, there can be additional, possibly inefficient, Nash equilibria, as the following example shows.

Example 21.13 Consider a system with R identical users with strictly concave utility function U . Suppose that user 1 chooses θ_1 so that $\bar{U}'(C; \theta_1) > U'(0)$, and every other user r chooses θ_r so that $\bar{U}'(0; \theta_r) < U'(C)$. Since $U'(C) \leq U'(0)$, it follows that (21.44) is satisfied for all users r . Thus this is a Nash equilibrium where the entire resource is allocated to user 1; however, the unique optimal solution to *SYSTEM* is symmetric, and allocates C/R units of the resource to each of the R users.

The equilibrium in the preceding example involves a “bluff”: user 1 declares such a high marginal utility at C that all other users concede. One way to preclude such equilibria is to enforce an assumption that guarantees participation. The next proposition assumes that all users have infinite marginal utility at zero allocation; this guarantees that all Nash equilibria are efficient.

Proposition 21.14 *Suppose that $U'_r(0) = \infty$ for all r . Suppose that θ is a Nash equilibrium. Then $\mathbf{d}(\theta)$ is an optimal solution to *SYSTEM*.*

PROOF Let $\mathbf{d} = \mathbf{d}(\theta)$. The proof follows by noting that all users must have positive allocations at equilibrium if $U'_r(0) = \infty$, from (21.44). Thus at equilibrium, for all users r, s we have $U'_r(d_r) = \bar{U}'(d_s; \theta_s)$. But this in turn implies that $U'_r(d_r) = U'_s(d_s)$ for all r, s , a sufficient condition for optimality for the problem *SYSTEM*. \square

Intuitively, for efficiency to hold, we need to have a number of actively “competing” users. In the previous result, this is guaranteed because every user will want strictly positive rate at any equilibrium.

The results of this section demonstrate that by relaxing the assumption that the resource allocation mechanism must set a single price, we can in fact significantly improve upon the efficiency guarantee of Theorem 21.10. It is critical to note that this gain in efficiency occurs only at *Nash equilibria*. The classical VCG mechanisms are unique in that they guarantee efficient outcomes as dominant strategy equilibria; it is straightforward to check that the SSVCG mechanisms described in this section will not have dominant strategy equilibria in general—e.g., the “bluff” example above is one such case.

21.5 Chapter Summary and Further Directions

This chapter considered the allocation of a single resource of fixed supply among multiple strategic users. We evaluated a variety of market mechanisms through Nash equilibria of the resulting resource allocation game. Our key insights are the following:

- (i) A simple proportional allocation mechanism, where each user receives a share of the resource in proportion to their bid, ensures full efficiency when users are price takers, and exhibits no worse than a 25% efficiency loss when users are price anticipators.
- (ii) In a natural class of mechanisms where users choose one-dimensional strategies, and the market sets a single price, the proportional allocation mechanism minimizes the worst-case efficiency loss when users are price anticipating; i.e., the best possible guarantee here is 75% of maximal aggregate utility.
- (iii) This guarantee can be improved if the mechanism is allowed to set one price per user. Using an adapted version of the VCG class of mechanisms, we can construct mechanisms that ensure fully efficient Nash equilibria.

Our investigation also reveals several further directions open for future research, including the following:

- (i) For the proportional allocation mechanism, we have proven a bound on the price of anarchy that shows that the ratio of the Nash equilibrium aggregate utility is no worse than $3/4$ the maximum possible aggregate utility. For nonatomic selfish routing (cf. Chapter 18), a similar price of anarchy result holds: the ratio of Nash cost to the optimal cost is no worse than $4/3$; furthermore, both proofs use the characterization of Nash equilibria as solutions to an optimization problem, with structure similar to the respective efficient optimization problems. These results are suggestive of perhaps a deeper generalization of price of anarchy for games with equilibria characterized as the solution to optimization problems.
- (ii) While Theorem 21.10 proves optimality of the proportional allocation mechanism in a reasonable class of mechanisms, the result depends critically on the assumption that all mechanisms in \mathcal{D} yield concave payoffs when agents are price anticipating. Given that some type of quasiconcavity assumption is typically necessary on payoffs to even guarantee existence of Nash equilibria, one might informally expect the result of Theorem 21.10 to hold even if Condition 2 is removed in the definition of \mathcal{D} . Whether this is in fact possible remains an open question.
- (iii) Our investigation shows, under reasonable assumptions, that with a single market-clearing price a 75% efficiency guarantee is possible, while with one price per user (the scalar strategy VCG approach), full efficiency is possible. This warrants further investigation: what is the exact trade-off between the number of prices and the efficiency guarantee possible? Furthermore, how does increasing the dimensionality of users' strategy affect this efficiency guarantee?

21.6 Notes

21.6.1 Section 20.2

Much of the material in this section is based on Chapter 2 of Johari (2004) and the corresponding paper (Johari and Tsitsiklis, 2004).

The mechanism discussed here was first studied in the context of communication networks by Kelly (1997). (See Chapter 22 for a discussion of the proportional allocation mechanism in congestion control algorithms for communication networks.) Theorem 21.1 is adapted from Kelly (1997), where it is proven in greater generality for an extension of the proportional allocation mechanism to a network context. This theorem is an extension of the classical *first fundamental theorem of welfare economics*; see Mas-Colell et al. (1995, Chapter 16), for details.

The first proof of uniqueness of Nash equilibrium for the proportional allocation mechanism was provided by La and Anantharam (2000). The most general result of existence and uniqueness, and the basis for the result in Theorem 21.2, is due to Hajek and Gopalakrishnan (2002); a less general result was proven by Maheswaran and Basar (2003). The explicit formulation of the problem *GAME* is given by Johari and Tsitsiklis (2004).

The price of anarchy result of Theorem 21.4 is due to Johari and Tsitsiklis (2004). The original proof of this result uses a two-step approach: it is first shown that the worst case is achieved using linear utility functions, and then the efficiency loss calculation is solved directly as a mathematical programming problem. The proof based on the problem *GAME* presented here is due to Roughgarden (2006), who also successfully applies the same method to efficiency loss calculations in several other games.

21.6.2 Section 20.3

Much of the material in this section is based on Chapter 5 of Johari (2004) and Section 4 of Johari and Tsitsiklis (2007).

The most closely related result to this section is presented by Maheswaran and Basar (2004). In their result, they consider mechanisms where each user r chooses a bid w_r , and the allocation is still made proportional to each player's bid. However, rather than assuming that every player pays w_r as in the standard proportional allocation mechanism, Maheswaran and Basar consider a class of mechanisms where the user pays $c(w_r)$, where c is a convex function. They show that in this class of mechanisms, the proportional allocation mechanism (i.e., a linear c) achieves the minimal worst-case efficiency loss when users are price anticipating.

Our work is substantially different, because we do not postulate that the mechanism must use the proportional rule (21.29) in allocating the resource; rather, this emerges as a consequence of rather simple assumptions on our mechanisms. We note that other works on inefficiency of resource allocation mechanisms, including Maheswaran and Basar (2004) and Yang and Hajek (2004), also assume a priori that allocations are made in proportion to users' bids.⁷ In this sense, our result lends a rigorous foundation to the

⁷ A notable exception is Sanghavi and Hajek (2004), which assumes that users pay their bid, and then designs an allocation rule to minimize worst case efficiency loss.

intuition that the proportional allocation rule (21.29) is a natural choice to determine the allocation among users.

21.6.3 Section 20.4

This section is based on Section 5.2 of the paper by Johari and Tsitsiklis (2007). Simultaneously and independently, a nearly identical formulation was developed by Yang and Hajek (2007). It is worth noting that Yang and Hajek and Maheswaran and Basar had earlier presented a resource allocation mechanism where users receive an allocation in proportion to their bids, but prices are chosen on an individualized basis (Maheswaran and Basar, 2004; Yang and Hajek, 2004); this mechanism can be seen to be a special case of the SSVCG mechanisms (Johari and Tsitsiklis, 2007).

Subsequent to the above work, several papers have presented related constructions of mechanisms that use limited communication yet achieve fully efficient Nash equilibria. Building on earlier work by Semret (1999), Dimakis et al. establish that a VCG-like mechanism where agents submit a pair (price and quantity requested) can achieve fully efficient equilibrium for a related resource allocation game (Dimakis et al., 2006). Stoescu and Ledyard consider the problem of resource allocation by building on the notion of minimal message spaces addressed in earlier literature on mechanism design, and build a class of efficient mechanisms with scalar strategy spaces (Stoescu and Ledyard, 2006).

The latter work of Stoescu and Ledyard recalls perhaps the most related reference (and most seminal) in this area by Reiter and Reichelstein (1988). Their paper calculates the minimal dimension of strategy space that would be necessary to achieve fully efficient Nash equilibria for a general class of economic models known as *exchange economies*. For our model, their bound evaluates to a strategy space per user of dimension $1 + 2/(R(R - 1))$, where R denotes the number of users. This is slightly higher than our result because Reiter and Reichelstein consider a much more general resource allocation problem.

Bibliography

- A. Dimakis, R. Jain, and J. Walrand. Mechanisms for efficient allocation in divisible capacity networks. *Proceedings of IEEE CDC*, pp. 1264–1269, 2006.
- B. Hajek and G. Gopalakrishnan. Do greedy autonomous systems make for a sensible Internet? Presented at the Conf. Stochastic Networks, Stanford University, 2002.
- R. Johari. *Efficiency Loss in Market Mechanisms for Resource Allocation*. PhD thesis, Massachusetts Institute of Technology, 2004.
- R. Johari and J.N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Math. Operat. Res.*, 29(3):407–435, 2004.
- R. Johari and J.N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. Management Science and Engineering Working Paper 07-04-6126-34, Stanford University, 2007.
- F.P. Kelly. Charging and rate control for elastic traffic. *Euro. Trans. Telecommun.*, 8:33–37, 1997.
- R.J. La and V. Anantharam. Charge-sensitive TCP and rate control in the Internet. In *Proc. IEEE INFOCOM*, pp. 1166–1175, 2000.

- R.T. Maheswaran and T. Basar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decis. Negotiation*, 12(5):361–395, 2003.
- R.T. Maheswaran and T. Basar. Social welfare of selfish agents: motivating efficiency for divisible resources. In *Proc. IEEE CDC*, pp. 1550–1555, 2004.
- A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, UK, 1995.
- S. Reichelstein and S. Reiter. Game forms with minimal message spaces. *Econometrica*, 56(3):661–692, 1988.
- T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proc. Intl. Congress of Mathematicians*, Vol. III, pp. 1071–1094, 2006.
- S. Sanghavi and B. Hajek. Optimal allocation of a divisible good to strategic buyers. In *Proc. IEEE CDC*, pp. 2748–2753, 2004.
- N. Semret. *Market Mechanisms for Network Resource Sharing*. PhD thesis, Columbia University, 1999.
- T.M. Stoescu and J. Ledyard. A pricing mechanism which implements a network rate allocation problem in Nash equilibria. 2006. Submitted.
- S. Yang and B. Hajek. An efficient mechanism for allocation of a divisible good and its application to network resource allocation. 2004. Preprint.
- S. Yang and B. Hajek. VCG-Kelly mechanisms for divisible goods: adapting VCG mechanisms to one-dimensional signals. To appear in *IEEE Journal on Selected Areas in Communications*, 2007.

Exercises

21.1 This exercise, together with the next one, studies the efficiency loss properties of the mechanisms defined in Example 21.9, by following the proof of Theorem 21.4. Suppose that $D(p, \theta) = \theta p^{-1/c}$, where $c \geq 1$. Suppose that given a utility system (C, R, \mathbf{U}) , a bid vector $\boldsymbol{\theta}$ is a Nash equilibrium, and let the resulting allocation vector be \mathbf{d} ; i.e., $d_r = D(p_D(\boldsymbol{\theta}), \theta_r)$.

- (a) Verify the Nash equilibrium conditions (21.31)–(21.32).
 (b) Show that \mathbf{d} is the unique solution to *GAME*, but where \hat{U}_r is defined as follows for each r :

$$\hat{U}_r(d_r) = \int_0^{d_r} \left(\frac{1 - z/C}{1 + (c-1)(z/C)} \right) U'_r(z) dz. \quad (\text{E1.1})$$

(Hint: rearrange the Nash equilibrium conditions (21.31)–(21.32).)

- (c) Show that \hat{U}_r satisfies Assumption 1.

21.2 Fix $D(p, \theta) = \theta p^{-1/c}$ and define \hat{U} as in the previous exercise. Define $\beta(D)$ according to (21.24), i.e.,

$$\beta(D) = \inf_{U \in \mathcal{U}} \inf_{C > 0} \inf_{0 \leq d, \bar{d} \leq C} \frac{U(d) + \hat{U}'(d)(\bar{d} - d)}{U(\bar{d})}.$$

- (a) Show that $\rho(D) \geq \beta(D)$. (Hint: first construct the variational inequality that identifies the optimality conditions for *GAME*, then argue as in the proof of Theorem 21.4.)
 (b) Show that $\beta(D) \geq G(c)$.

- (c) Using a construction analogous to the proof of Theorem 21.4, show that for any δ there exists a utility system for which the ratio of Nash aggregate utility to the maximum aggregate utility is no more than $G(c) + \delta$. Conclude that $\rho(D) = G(c)$.
- 21.3** Show by example that a VCG mechanism does not necessarily charge each user the same per-unit price for the resource.