Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games*

Tim Roughgarden† Éva Tardos‡

September 24, 2003

Abstract

Equilibria in noncooperative games are typically inefficient, as illustrated by the Prisoner's Dilemma. In this paper, we quantify this inefficiency by comparing the payoffs of equilibria to the payoffs of a "best possible" outcome. We study a nonatomic version of the congestion games defined by Rosenthal (1973a), and identify games in which equilibria are *approximately optimal* in the sense that no other outcome achieves a significantly larger total payoff to the players—games in which optimization by individuals approximately optimizes the social good, in spite of the lack of coordination between players. Our results extend previous work on traffic routing games.

*We thank Lou Billera for encouraging us to explore generalizations of our previous work on traffic routing, Robert Rosenthal for introducing us to his congestion games, and the anonymous referee for helpful comments and corrections. The first author thanks Amir Ronen for comments on a previous paper (Roughgarden, 2002a) that facilitated a generalization of the results therein to the nonatomic congestion games of this paper.

†Department of Computer Science, Cornell University, Ithaca NY 14853. Supported by an NSF Fellowship, a Cornell University Fellowship, and ONR grant N00014-98-1-0589. Email: `timr@cs.cornell.edu`.

‡Department of Computer Science, Cornell University, Ithaca NY 14853. Supported by NSF grant CCR-9700163 and ONR grant N00014-98-1-0589. Email: `eva@cs.cornell.edu`.

1 Introduction

Equilibria in noncooperative games are typically inefficient, as illustrated by the classical Prisoner's Dilemma (Rapoport and Chammah, 1965). In this paper, we aim to *quantify* this inefficiency by comparing the payoffs of equilibria to the payoffs of a "best possible" outcome. To make such a comparison and to define a notion of optimality, we assume that aggregating the payoffs of different players is meaningful; it is then natural to study the sum of player payoffs of noncooperative equilibria and of coordinated outcomes. As a simple example, the total number of years of imprisonment handed down to the prisoners in the Nash equilibrium of the Prisoner's Dilemma (with both prisoners defecting) can be compared to the same quantity in the (optimal) coordinated outcome. Our goal is to identify natural classes of games in which equilibria are guaranteed to be *approximately optimal*, in the sense that no other outcome achieves a significantly larger sum of player payoffs; these are games in which optimization by individuals approximately optimizes the social good, in spite of the competition and lack of coordination between players.

We study a nonatomic version (in the sense of Schmeidler (1973)) of the congestion games defined by Rosenthal (1973a). Congestion games are noncooperative games in which the payoff to each player depends only on the player's strategy and on the number of other players choosing the same or some "interfering" strategy. Congestion games have been used to model traffic behavior in road and communication networks, competition among firms for production processes, and migration of animals between different habitats (Milchtaich, 1996; Quint and Shubik, 1994; Rosenthal, 1973a; Roughgarden and Tardos, 2002). Rosenthal (1973a) studied games with a finite number of discrete players; more recently, several authors have considered nonatomic congestion games in which there is a continuum of players (Blonski, 1999; Friedman, 1996; Milchtaich, 1996, 2001).

The idea of quantifying the inefficiency of noncooperative equilibria using notions of approximation is due to Koutsoupias and Papadimitriou (1999), who proved upper and lower bounds on the worst-possible objective function value of a Nash equilibrium relative to that of an optimal solution in a simple load-balancing game. The authors, motivated by networking issues, subsequently applied this idea to a classical model of traffic routing with noncooperative players (Wardrop, 1952) and determined the worst-case inefficiency of equilibria in this model (Roughgarden, 2002a,b; Roughgarden and Tardos, 2002).

In this paper, we extend techniques developed for analyzing traffic equilibria (Roughgarden, 2002a; Roughgarden and Tardos, 2002) to a large class of nonatomic congestion games. The work presented here includes all previous bounds on the inefficiency of traffic equilibria as special cases, and demonstrates that traffic routing games in no way exhaust the supply of games with provably near-optimal equilibria. Our results also suggest that studying the degree to which equilibria approximate natural objective functions may be tractable and enlightening in many other game-theoretic contexts.

Organization and Overview

In Section 2 we define our model, study a simple example, and state several lemmas about equilibria in nonatomic congestion games.

In Sections 3 and 4 we quantify the inefficiency of equilibria by investigating the worst possible ratio between the total cost incurred by players in an equilibrium and in an outcome of minimum-possible total cost (it is technically convenient to assign costs, rather than payoffs, to outcomes). This ratio has recently been dubbed "the price of anarchy" by Papadimitriou (2001). Section 3 is devoted to a "quick and dirty" upper bound on the worst-case inefficiency of equilibria in certain nonatomic congestion games; this method is easy to apply but does not always give a best possible bound. In Section 4 we give, by more sophisticated methods, an exact bound on this worst-case inefficiency under mild conditions on the cost functions of the game. We also show in this section that simple games always furnish worst-possible examples for the inefficiency of equilibria.

In Section 5, we consider a more general class of nonatomic congestion games, where the cost incurred by players in an equilibrium can be arbitrarily larger than that in other outcomes. We show that a meaningful upper bound on the inefficiency of equilibria is nevertheless possible in this general setting; specifically, we prove that the cost of an equilibrium is bounded above by that of the best possible outcome between an increased number of players.

2 Preliminaries

We consider a finite ground set E of *elements*, with each element e possessing a *cost function* $c_e(\xi)$. We assume that each cost function is nonnegative and is continuous and nondecreasing in its argument. There are k *player types* $1, 2, \ldots, k$, and for each player type i there is a finite multiset $\mathcal{S}_i$ of subsets of E, called the *strategy set* of players of type i. Elements of $\mathcal{S}_i$ are called *strategies*. The continuum of players of type i is represented by the interval $[0, n_i]$ endowed with Lebesgue measure. To a player type i, a strategy $S \in \mathcal{S}_i$, and an element $e \in S$, we associate a positive *rate of consumption* $a_{S,e}$ that defines the amount of congestion contributed to element e by players of type i selecting strategy S. A *nonatomic congestion game (NCG)* is defined by a 5-tuple $(E, c, \mathcal{S}, n, a)$.

Example 2.1 A classical and intensively studied type of NCG is the traffic routing model of Wardrop (1952). Here, the set E is the edge set of a directed graph, player types correspond to different origin-destination vertex pairs, strategy sets are sets of directed origin-destination paths, and the cost function c_e describes the delay experienced by traffic traversing edge e as a function of the edge congestion (typically, all rates of consumption are 1).

By an *action distribution*, we mean a vector x of nonnegative reals with components indexed by the disjoint union of $\mathcal{S}_1, \ldots, \mathcal{S}_k$, with the property that $\sum_{S \in \mathcal{S}_i} x_S = n_i$ for each player type i. We interpret x_S as the measure of the set of players selecting strategy S. We abuse notation and write x_e for the total amount of congestion induced on element e by the action distribution x:

$$x_e = \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} a_{S,e} x_S.$$

We next define the disutility experienced by players with respect to an action distribution. With respect to an action distribution x, players of type i selecting strategy $S \in \mathcal{S}_i$ incur a *cost* $c_S(x)$ defined by

$$c_S(x) = \sum_{e \in S} a_{S,e} c_e(x_e).$$

The cost incurred by a player is thus the sum of the costs incurred on the elements of its strategy, each of which in turn is the base cost $c_e(x_e)$ of the element e (with respect to the congestion induced by action distribution x) times the rate $a_{S,e}$ at which the player consumes the element. We emphasize that the cost incurred by a player depends only on its type, its chosen strategy, and on the congestion experienced by the elements in its strategy—the identities of the players that induce the congestion are of no importance.

We next introduce an objective function to measure how "good" an action distribution is. Define the *social cost* $C(x)$ of an action distribution x as the total disutility experienced by the players, so that

$$C(x) = \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} c_S(x) x_S.$$

We will be interested in action distributions with smallest possible social cost, and will call such an action distribution *optimal*. Any nonatomic congestion game admits an optimal action distribution because the social cost function C is assumed continuous and the space of all action distributions (a subset of $\mathcal{R}^{\mathcal{S}_1} \times \cdots \times \mathcal{R}^{\mathcal{S}_k}$) is compact.

Remark 2.2 Nonatomic games are traditionally presented in a more general way, as follows. The basic notion is that of a *strategy profile*, defined as a (Lebesgue) measurable function σ from a compact interval I to a finite-dimensional simplex of mixed strategies (Schmeidler, 1973). A strategy profile then naturally induces an action distribution x, with x_S obtained by integrating the coordinate of σ corresponding to strategy S over the interval I (with respect to Lebesgue measure). Conversely, every action distribution is induced by some strategy profile. When all players select pure strategies, passing from strategy profiles to action distributions can be viewed as aggregating players according to their chosen strategies and ignoring their identities. We will not directly study strategy profiles in this paper, because our investigation of nonatomic games concerns only a quantity (the social cost) that, by definition, depends only on the action distribution of a strategy profile. We will work only with the elementary notion of an action distribution in what follows.

We next define our notion of equilibrium in NCGs. Informally, equilibria are the action distributions induced by strategy profiles (see Remark 2.2) in which all players outside a set of measure zero play best response strategies.

Definition 2.3 An action distribution x for a nonatomic congestion game $(E, c, \mathcal{S}, n, a)$ is an *equilibrium* if for each player type $i = 1, 2, \ldots, k$ and strategies $S_1, S_2 \in \mathcal{S}_i$ with $x_{S_1} > 0$, $c_{S_1}(x) \leq c_{S_2}(x)$.

Example 2.4 Define a nonatomic congestion game $\Gamma_p = (E, c, \mathcal{S}, n, a)$ for a positive integer p as follows. Let $E = \{1, 2\}$ be a ground set with two elements and cost functions $c_1(\xi) = 1$ and $c_2(\xi) = \xi^p$. There is one player type with $n_1 = 1$, and two singleton strategies $\mathcal{S}_1 = \{\{1\}, \{2\}\}$ with unit rates of consumption $a_{\{1\},1} = a_{\{2\},2} = 1$. For any p, the action distribution $x_{\{1\}} = 0$, $x_{\{2\}} = 1$ is an equilibrium for Γ_p with social cost 1. A superior (indeed, optimal) action distribution is $x_{\{1\}} = 1 - (p+1)^{-1/p}$, $x_{\{2\}} = (p+1)^{-1/p}$, which has social cost $1 - p \cdot (p+1)^{-(p+1)/p} < 1$. This quantity is 0.75 when $p = 1$, is approximately 0.615 when $p = 2$ and 0.528 when $p = 3$, and tends to 0 as $p \to \infty$.

Example 2.4 demonstrates that, in an arbitrary NCG, an equilibrium may have social cost arbitrarily larger than that of other action distributions. This fact motivates our work identifying classes of NCGs in which equilibria are not arbitrarily inefficient (Sections 3 and 4) and seeking approaches to bounding the social cost of equilibria in general NCGs that do not directly compare the social cost of optimal and equilibrium action distributions of the same game (Section 5).

Equilibrium action distributions in nonatomic congestion games with continuous, nondecreasing cost functions enjoy several nice properties. We now state three of them.

Proposition 2.5 *Every nonatomic congestion game admits an equilibrium action distribution.*

Proposition 2.5 is a special case of very general theorems about the existence of equilibria in nonatomic games (Mas-Colell, 1984; Rath, 1992; Schmeidler, 1973) and can also be proved by a straightforward generalization of techniques used to prove existence of traffic equilibria (Beckmann *et al.*, 1956; Dafermos and Sparrow, 1969).

Proposition 2.6 *Distinct equilibria for a nonatomic congestion game have equal social cost.*

Proposition 2.6 follows from work of Milchtaich (2000) or from an easy generalization of an analogous result for traffic equilibria (Beckmann *et al.*, 1956; Dafermos and Sparrow, 1969).

Proposition 2.7 *Let x be an equilibrium action distribution for the nonatomic congestion game $(E, c, \mathcal{S}, n, a)$. For each player type i, there is a real number $c_i(x)$ such that every strategy $S \in \mathcal{S}_i$ with $x_S > 0$ satisfies $c_S(x) = c_i(x)$. The social cost of x is then*

$$C(x) = \sum_{i=1}^{k} c_i(x) n_i.$$

The first assertion in Proposition 2.7 is immediate from the definition of an equilibrium; the second is immediate from the first and the definition of social cost.

Our final preliminary result is a simple but useful equivalent definition of social cost, that will account for the social cost of an action distribution on an element-by-element (rather than strategy-by-strategy) basis.

Proposition 2.8 *The social cost $C(x)$ of an action distribution x for the nonatomic congestion game $(E, c, \mathcal{S}, n, a)$ satisfies*

$$C(x) = \sum_{e \in E} c_e(x_e) x_e.$$

Proof: Expanding the definitions and rearranging, we obtain

$$C(x) = \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} x_S \sum_{e \in S} a_{S,e} c_e(x_e) = \sum_{e \in E} c_e(x_e) \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i : e \in S} a_{S,e} x_S = \sum_{e \in E} c_e(x_e) x_e.$$

■

3 A "Quick and Dirty" Upper Bound

In this section and the next, we demonstrate that nontrivial upper bounds on the inefficiency of equilibria in nonatomic congestion games are possible. In this section we present a simple method that gives a nontrivial but suboptimal bound for a large class of NCGs; in the next, we pursue exact worst-case bounds by more sophisticated methods.

For a NCG Γ, define $\rho(\Gamma)$ by $C(x)/C(x^*)$ where x is an equilibrium and x^* is an optimal action distribution for Γ; this ratio is well defined by Propositions 2.5 and 2.6 except in the degenerate case where $C(x^*) = 0$. In this case, x^* is also an equilibrium and hence (by Proposition 2.6) $C(x) = 0$; we then define $\rho(\Gamma) = 1$.

We next give a method for bounding the ρ-value of NCGs with cost functions that are, in some sense, "not too steep". For example, we will see that NCGs with cost functions that are bounded-degree polynomials with nonnegative coefficients have bounded ρ-value (independent of E, $\mathcal{S}$, n and a).

Theorem 3.1 *Suppose the nonatomic congestion game* $\Gamma = (E, c, \mathcal{S}, n, a)$ *and the constant* $\eta \geq 1$ *satisfy*

$$\xi \cdot c_e(\xi) \leq \eta \cdot \int_0^{\xi} c_e(t)dt$$

for all elements e *and all positive real numbers* ξ. *Then*

$$\rho(\Gamma) \leq \eta.$$

Proof: The proof hinges on the following fact: the equilibria for Γ are precisely the optimal action distributions for the NCG $\tilde{\Gamma} = (E, \tilde{c}, \mathcal{S}, n, a)$, where $\tilde{c}_e(0) = c_e(0)$ and $\tilde{c}_e(\xi) = [\int_0^{\xi} c_e(t)dt]/\xi$ for $\xi > 0$. The proof of this fact is a routine generalization of an analogous result for traffic equilibria (Beckmann *et al.*, 1956; Dafermos and Sparrow, 1969); the idea is to model the optimization problem of computing the optimal action distribution for $\tilde{\Gamma}$ as a convex nonlinear program, and to check that the Karush-Kuhn-Tucker conditions (Peressini *et al.*, 1988), which provide necessary and sufficient conditions for optimality, are equivalent to the equilibrium conditions of Definition 2.3 for the game Γ. We remark that Rosenthal (1973a; 1973b) showed that this fact and its proof also have analogues for atomic congestion games. In turn, his work motivated the recent study of *potential games*—games more general than congestion games, in which Nash equilibria can be interpreted as maximizing an appropriate objective function, called a *potential function*. We refer the reader to the seminal paper of Monderer and Shapley (1996) and the survey of Voorneveld *et al.* (1999) for a detailed discussion of these topics.

We now show that an equilibrium for Γ optimizes an objective function that is at most a factor of η away from the social cost $C(\cdot)$. Let x and x^* denote an equilibrium and an optimal action distribution for Γ, respectively. We then have

$$\begin{aligned} C(x) &= \sum_{e \in E} c_e(x_e) x_e && (1) \\ &\leq \eta \sum_{e \in E} \int_0^{x_e} c_e(t)dt && (2) \\ &\leq \eta \sum_{e \in E} \int_0^{x_e^*} c_e(t)dt && (3) \\ &\leq \eta \sum_{e \in E} c_e(x_e^*) x_e^* && (4) \\ &= \eta \cdot C(x^*) && (5) \end{aligned}$$

where (1) and (5) follow from Proposition 2.8, (2) from the hypothesis, (3) from the fact that the equilibrium x minimizes the function $\sum_e \int_0^{\xi_e} c_e(t)dt$, and (4) from the assumption that every cost function c_e is nondecreasing. ■

For example, Theorem 3.1 yields the following corollary for NCGs with cost functions that are polynomials with nonnegative coefficients.

Corollary 3.2 *Suppose every cost function of the NCG Γ is a polynomial with nonnegative coefficients and degree at most p. Then,*

$$\rho(\Gamma) \leq p + 1.$$

4 An Exact Bound for Most Classes of Cost Functions

In this section we undertake a deeper study of the inefficiency of equilibria in nonatomic congestion games. We saw in Example 2.4 that for arbitrary (continuous, nondecreasing) cost functions, equilibria can be arbitrarily inefficient. On the other hand, in the previous section we discovered that restricting the class of allowable cost functions (to degree-bounded polynomials, say) permits a universal upper bound on this inefficiency, independent of all other parameters (such as the number of elements and the number of player types). These two observations motivate our current ambition: to precisely compute the worst-case ratio in social cost between an equilibrium and an optimal action distribution with respect to any (fixed but arbitrary) class of allowable cost functions. We will realize this ambition subject to two mild assumptions on the allowable cost functions. We record these assumptions next.

Definition 4.1 A class $\mathcal{C}$ of cost functions is *homogeneous* if $c(0) = 0$ for all $c \in \mathcal{C}$ and *inhomogeneous* otherwise.

Definition 4.2 A cost function c is *standard* if it is differentiable and if the function $\xi \cdot c(\xi)$ is convex on $[0, \infty)$. A class $\mathcal{C}$ of cost functions is *standard* if it contains only standard cost functions and includes at least one nonzero function.

Convex differentiable cost functions are standard, as are some nonconvex functions such as $\log(1+\xi)$. Approximations of step functions are examples of nonstandard cost functions.

Throughout this section, we restrict attention to NCGs with cost functions drawn from an inhomogeneous standard class.[1] We now proceed in two steps. We begin with a lower bound on the worst-case ratio between the social cost of an equilibrium and of an optimal action distribution, and then prove a matching upper bound.

4.1 Lower Bounding the Inefficiency of Equilibria

We first give a fairly obvious lower bound on the worst-case value of ρ with respect to NCGs with cost functions drawn from a standard class $\mathcal{C}$ containing the constant functions (cost functions of the form $c(\xi) = \beta$ for a scalar $\beta > 0$). We will then show that this lower bound is in fact valid with respect to any inhomogeneous standard class of cost functions.

Let $\mathcal{C}$ be a standard class of cost functions containing all of the constant functions. We will construct a bad example (a NCG with cost functions in $\mathcal{C}$ and an inefficient equilibrium) by mimicking Example 2.4. Define elements E, strategy set $\mathcal{S}_1$, and rates a as in Example 2.4. We assign an arbitrary nonzero cost function $c_2 \in \mathcal{C}$ to element 2 and pick an arbitrary positive value for n_1, large enough so that $c_2(n_1) > 0$. We conclude our definition of the game $\Gamma = (E, c, \mathcal{S}, n, a)$ by assigning element 1 the constant cost function $c_1(\xi) = c_2(n_1)$, which by hypothesis lies in $\mathcal{C}$.

The action distribution $x_{\{1\}} = 0$, $x_{\{2\}} = n_1$ is an equilibrium for Γ with social cost $c_2(n_1)n_1$. To compute the optimal action distribution, we define the *marginal social cost function* c_e^* corresponding to the cost function c_e as the derivative of the social cost $\xi \cdot c_e(\xi)$ incurred on edge e. We note that $c_e^*(\xi) = c_e(\xi) + \xi \cdot c_e'(\xi)$; since cost functions are assumed nondecreasing, $c_e^* \geq c_e$ on $[0, \infty)$. Returning to the game Γ, we define the action distribution x^* so that marginal costs are equalized ($c_1^*(x_1^*) = c_2^*(x_2^*)$); this is possible since c_1^* is everywhere equal to $c_2(n_1)$, $c_2^*(0) = c_2(0)$, and $c_2^* \geq c_2$ with c_2^* continuous. It is easy to check that, since the cost functions c_1 and c_2 are standard, x^* is an optimal action distribution for Γ.[2] The social cost of x^* is $[\lambda\mu + (1-\lambda)]c_2(n_1)n_1$, where $\lambda = x_2^*/x_2$ and $\mu = c_2(x_2^*)/c_2(x_2)$, so $\rho(\Gamma) = [\lambda\mu + (1-\lambda)]^{-1}$.

Motivated by this example, we next assign to each standard class $\mathcal{C}$ of cost functions a number $\alpha(\mathcal{C}) \in [1, \infty]$ that, when $\mathcal{C}$ contains the constant functions, lower bounds the worst-case ρ-value occurring in NCGs with cost functions in $\mathcal{C}$.

Definition 4.3 Let $\mathcal{C}$ be a standard class of cost functions. For a nonzero cost function $c \in \mathcal{C}$, we define $\alpha(c)$ by

$$\alpha(c) = \sup_{n>0\,:\,c(n)>0} [\lambda\mu + (1-\lambda)]^{-1}$$

[1] The differentiability condition of Definition 4.2 is for simplicity of presentation and can be relaxed.

[2] More generally, the optimal action distributions for a NCG $(E, c, \mathcal{S}, n, a)$ with standard cost functions c are precisely the equilibria for the NCG $(E, c^*, \mathcal{S}, n, a)$. This fact follows from the equivalence between equilibrium conditions for NCGs and optimality conditions for certain convex programs described in the proof of Theorem 3.1; as discussed in that proof, this is tantamount to asserting that NCGs are (nonatomic) potential games in the sense of Monderer and Shapley (1996). It also explains our notational convention of writing x^* for an optimal action distribution and c^* for a marginal social cost function: the functions c^* are "optimal cost functions" in that the optimal action distributions arise as equilibria with respect to c^*.

where $\lambda \in [0,1]$ satisfies $c^*(\lambda n) = c(n)$ and $\mu = c(\lambda n)/c(n) \in [0,1]$. We define $\alpha(\mathcal{C})$ by

$$\alpha(\mathcal{C}) = \sup_{0 \neq c \in \mathcal{C}} \alpha(c).$$

Existence of the scalar λ follows from an argument already given above. For most cost functions λ is uniquely determined by c and n; otherwise, our assumption that each cost function c is standard ensures that $\alpha(c)$ is well defined (i.e., that $[\lambda\mu + (1-\lambda)]^{-1}$ is independent of the choice of λ satisfying $c^*(\lambda n) = c(n)$).

Remark 4.4 The expression for $\alpha(\mathcal{C})$ may not look easy to work with, but it simplifies considerably in many cases of interest. For example, if $\mathcal{C}$ is the set of polynomials with degree at most p and nonnegative coefficients, then $\alpha(\mathcal{C}) = [1 - p \cdot (p+1)^{-(p+1)/p}]^{-1}$; see Roughgarden (2002b) for a derivation of this fact and for further examples motivated by applications in networking and queueing theory.

It is immediate that if $\mathcal{C}$ is a standard class of cost functions containing the constant functions, then there are NCGs with ratio ρ arbitrarily close to $\alpha(\mathcal{C})$. Our next theorem extends this statement to inhomogeneous standard classes of cost functions.

Theorem 4.5 *Let $\mathcal{C}$ denote an inhomogeneous standard class of cost functions. Let $\mathcal{G}$ denote the set of nonatomic congestion games with cost functions in $\mathcal{C}$. Then*

$$\sup_{\Gamma \in \mathcal{G}} \rho(\Gamma) \geq \alpha(\mathcal{C}).$$

The idea of the proof of Theorem 4.5 is that, for a class $\mathcal{C}$ not containing the constant functions, elements with constant cost function can nonetheless be "simulated" with multiple elements all endowed with the same cost function from $\mathcal{C}$. The details are not difficult, but are somewhat technical and are therefore deferred to the Appendix.

Remark 4.6 Theorem 4.5 does not hold if the inhomogeneity hypothesis is omitted. Indeed, if $\mathcal{C}$ is the homogeneous class comprising the monomials with a nonnegative coefficient and degree exactly p for some $p > 0$, then $\alpha(\mathcal{C}) > 1$ but $\rho(\Gamma) = 1$ for all NCGs Γ with cost functions in $\mathcal{C}$ (see, for example, Dafermos and Sparrow (1969) or Barro and Romer (1987)).

4.2 Upper Bounding the Inefficiency of Equilibria

We now prove the following: for any nonatomic congestion game Γ with cost functions in the standard class $\mathcal{C}$, $\rho(\Gamma) \leq \alpha(\mathcal{C})$. This result provides a matching upper bound to Theorem 4.5 and completely resolves the worst-case inefficiency of equilibria of NCGs with respect to an inhomogeneous standard class of cost functions.

Theorem 4.7 *If Γ is a nonatomic congestion game with cost functions in the standard class $\mathcal{C}$, then then $\rho(\Gamma) \leq \alpha(\mathcal{C})$.*

Proof: Let x^* and x be optimal and equilibrium action distributions, respectively, for a nonatomic congestion game $\Gamma = (E, c, \mathcal{S}, n, a)$ with cost functions in the standard class $\mathcal{C}$. We begin by decomposing the social cost of an optimal action distribution—the quantity that we wish to lower bound—into two expressions more closely related to the social cost of an equilibrium.

Our assumption that each cost function c_e is standard ensures that each marginal social cost function c_e^* is nondecreasing. We can therefore lower bound the social cost $c_e(x_e^*)x_e^*$ incurred on edge e using a linear approximation of the function $c_e(\xi)\xi$ at the point $\lambda_e x_e$, where $\lambda_e \in [0,1]$ solves $c_e^*(\lambda_e x_e) = c_e(x_e)$:

$$\begin{aligned} c_e(x_e^*)x_e^* &= c_e(\lambda_e x_e)\lambda_e x_e + \int_{\lambda_e x_e}^{x_e^*} c_e^*(x)dx \\ &\geq c_e(\lambda_e x_e)\lambda_e x_e + (x_e^* - \lambda_e x_e)c_e^*(\lambda_e x_e) \\ &= c_e(\lambda_e x_e)\lambda_e x_e + (x_e^* - \lambda_e x_e)c_e(x_e). \end{aligned}$$

Summing this inequality over all elements and applying Proposition 2.8 to the left-hand side, we obtain

$$C(x^*) \geq \sum_e [c_e(\lambda_e x_e)\lambda_e x_e + (x_e^* - \lambda_e x_e)c_e(x_e)].$$

We next rewrite this expression so that it enjoys a close connection with the value $\alpha(\mathcal{C})$ of Definition 4.3—the upper bound we wish to prove for $\rho(\Gamma)$. Precisely, we can write

$$C(x^*) \geq \sum_e [\mu_e \lambda_e x_e + (1-\lambda_e)x_e]c_e(x_e) + \sum_e [x_e^* - x_e]c_e(x_e)$$

where, following Definition 4.3, μ_e is defined as $c_e(\lambda_e x_e)/c_e(x_e)$ (if $c_e(x_e) = 0$, put $\mu_e = 1$). This first sum of this expression is, as desired, closely related to $\alpha(\mathcal{C})$; the second sum can be regarded as an "error term".

We claim that this error term is nonnegative and can therefore be dropped without harm. To prove it, we first use the method of proof of Proposition 2.8 to write

$$\sum_e c_e(x_e)x_e = \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} c_S(x)x_S$$

and

$$\sum_e c_e(x_e)x_e^* = \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} c_S(x)x_S^*.$$

By Proposition 2.7, the first sum equals $\sum_{i=1}^{k} c_i(x)n_i$, where $c_i(x)$ is the common cost (with respect to x) of all strategies S in $\mathcal{S}_i$ with $x_S > 0$. By Definition 2.3, $c_S(x) \geq c_i(x)$ for all $S \in \mathcal{S}_i$; the second sum is thus bounded below by $\sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} c_i(x)x_S^* = \sum_{i=1}^{k} c_i(x)n_i$. It follows that $\sum_e c_e(x_e)x_e \leq \sum_e c_e(x_e)x_e^*$ and hence the error term $\sum_e [x_e^* - x_e]c_e(x_e)$ is nonnegative, as claimed.

We have now established the inequality

$$C(x^*) \geq \sum_e [\mu_e \lambda_e x_e + (1-\lambda_e)x_e]c_e(x_e).$$

By Definition 4.3, $\mu_e\lambda_e + (1-\lambda_e) \geq 1/\alpha(\mathcal{C})$ for each element e; thus, the quantities $[\mu_e\lambda_e x_e + (1-\lambda_e)x_e]c_e(x_e)$ and $c_e(x_e)x_e$ differ by at most an $\alpha(\mathcal{C})$ factor for each element e. Summing over all elements and applying Proposition 2.8, we find that the social costs of x^* and x also differ by at most an $\alpha(\mathcal{C})$ factor:

$$C(x^*) \geq \frac{1}{\alpha(\mathcal{C})} \sum_e c_e(x_e)x_e = \frac{C(x)}{\alpha(\mathcal{C})}.$$

The theorem is proved. ■

For example, if $\mathcal{C}_p$ is the class of polynomials with degree at most p and nonnegative coefficients, then the worst-case ρ-value of NCGs with cost functions in $\mathcal{C}_p$ is precisely $[1 - p \cdot (p+1)^{-(p+1)/p}]^{-1}$. This worst-case bound is asymptotically $\Theta(p/\ln p)$ as $p \to \infty$ and is realized (for all p) by the simple NCGs of Example 2.4.

In fact, Theorem 4.7 and the proof of Theorem 4.5 (given in the Appendix) show more generally that simple games always furnish worst-possible examples of the inefficiency of equilibria. This is immediate from Definition 4.3 and Theorem 4.7 when the class of allowable cost functions is standard and includes all of the constant functions, with two-element NCGs analogous to Example 2.4 providing worst-case examples. The proof of Theorem 4.5 shows how an element with constant cost function can be "simulated" with many elements with nonconstant cost functions; it follows from this construction that NCGs with only one player type, disjoint singleton strategies, and unit rates of consumption are worst-case examples with respect to a standard class $\mathcal{C}$ of allowable cost functions satisfying $\{c(0) : c \in \mathcal{C}\} \supseteq (0, \infty)$. A similar simulation argument (given in the proof of Theorem 4.5) shows that, under the weaker assumption that the class of allowable cost functions is inhomogeneous, NCGs with one player type, disjoint strategies, and unit rates of consumption give worst-possible ρ-values.

Corollary 4.8 *Let $\mathcal{C}$ be an inhomogeneous standard class of cost functions. Let $\mathcal{G}$ denote the set of NCGs with cost functions in $\mathcal{C}$ and $\mathcal{G}' \subseteq \mathcal{G}$ the NCGs with one player type, mutually disjoint strategies, and $a_{S,e} = 1$ for all strategies S and elements e. Let $\mathcal{G}'_m \subseteq \mathcal{G}'$ denote the NCGs that, in addition, possess only m singleton strategies. Then*

$$\sup_{\Gamma \in \mathcal{G}'} \rho(\Gamma) = \sup_{\Gamma \in \mathcal{G}} \rho(\Gamma).$$

If for each positive scalar $\beta > 0$ there is a cost function $c \in \mathcal{C}$ with $c(0) = \beta$, then

$$\sup_{\Gamma \in \cup_m \mathcal{G}'_m} \rho(\Gamma) = \sup_{\Gamma \in \mathcal{G}} \rho(\Gamma).$$

If $\mathcal{C}$ contains the constant functions, then

$$\sup_{\Gamma \in \mathcal{G}'_2} \rho(\Gamma) = \sup_{\Gamma \in \mathcal{G}} \rho(\Gamma).$$

5 A Weaker Type of Guarantee for Arbitrary Cost Functions

In Example 2.4, we observed that the ratio $\rho(\Gamma)$ can be arbitrarily large for NCGs Γ with arbitrary cost functions. In the previous two sections, we evaded this difficulty by imposing additional restrictions (beyond continuity and monotonicity) on the class of allowable cost functions. On the other hand, upper-bounding the worst-possible value of the ratio ρ is not the only available method to prove limits on the inefficiency of equilibria. In this section, we avoid making additional assumptions on element cost functions by pursuing a different, weaker type of guarantee on the worst-possible inefficiency of equilibria. Rather than upper bounding the social cost of an equilibrium relative to that of an optimal action distribution (an impossible feat in the current general setting), we upper bound this social cost by that of an optimal action distribution for the same game with *twice as many players* for each player type. This type of guarantee is nontrivial (it is not difficult to find natural classes of games for which this bound fails), and can also have a natural meaning in concrete scenarios. For example, this bound has the following interpretation for routing traffic in certain networks where each edge possesses a capacity (or bandwidth): the social cost of a traffic equilibrium after doubling the capacity of the network is bounded above by the cost of the best coordinated outcome in the original network (Roughgarden and Tardos, 2002). Thus, the losses due to noncooperative behavior in such a network can be offset with a moderate investment in network hardware.

We now state and prove our result bounding the social cost of an equilibrium by the social cost of an optimal action distribution in a game with an increased number of players.

Theorem 5.1 *If x is an equilibrium for $(E, c, \mathcal{S}, n, a)$ and x^* is an action distribution for $(E, c, \mathcal{S}, 2n, a)$, then $C(x) \leq C(x^*)$.*

Proof: Let x, x^* satisfy the hypotheses of the theorem. For each player type $i = 1, \ldots, k$, let $c_i(x)$ be the common cost of every strategy $S \in \mathcal{S}_i$ with $x_S > 0$ and write $C(x) = \sum_{i=1}^{k} c_i(x)n_i$ (see Proposition 2.7). We will define new cost functions $\bar{c}$ that both approximate the original ones (in the sense that the social cost of an action distribution with respect to cost functions $\bar{c}$ is close to its original social cost) and allow us to easily lower bound the cost (with respect to $\bar{c}$) of x^*. Precisely, we define new cost functions $\bar{c}$ as follows:

$$\bar{c}_e(\xi) = \begin{cases} c_e(x_e) & \text{if } \xi \leq x_e \\ c_e(\xi) & \text{if } \xi \geq x_e. \end{cases}$$

First we show that the new cost functions $\bar{c}$ approximate the original ones. For any element e, $\bar{c}_e(\xi) - c_e(\xi)$ is zero for $\xi \geq x_e$ and bounded above by $c_e(x_e)$ for $\xi < x_e$, so $\xi(\bar{c}_e(\xi) - c_e(\xi)) \leq c_e(x_e)x_e$ for all $\xi \geq 0$. Using this fact and Proposition 2.8, we find that evaluating x^* with cost functions $\bar{c}$ (rather than c) increases its social cost by at most an additive $C(x)$ factor:

$$\sum_{e \in E} \bar{c}_e(x_e^*)x_e^* - \sum_{e \in E} c_e(x_e^*)x_e^* = \sum_{e \in E} x_e^*(\bar{c}_e(x_e^*) - c_e(x_e^*)) \leq \sum_{e \in E} c_e(x_e)x_e = C(x).$$

To complete the proof, we will show that the social cost of the action distribution x^* with respect to the cost functions $\bar{c}$ is lower bounded by $2C(x)$. Since $\bar{c}_e(\xi) \geq c_e(x_e)$ for every element $e \in E$ and nonnegative real ξ, by definition of the values $c_i(x)$ we have $\bar{c}_S(x^*) \geq c_i(x)$ for every player type i and strategy $S \in \mathcal{S}_i$. The cost of x^* with respect to $\bar{c}$ can therefore be bounded below as follows:

$$\sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} \bar{c}_S(x^*) x_S^* \geq \sum_{i=1}^{k} \sum_{S \in \mathcal{S}_i} c_i(x) x_S^* = \sum_{i=1}^{k} 2c_i(x) n_i = 2C(x).$$

■

The same proof shows that, more generally, if x is an equilibrium for $(E, c, \mathcal{S}, n, a)$ and x^* is an action distribution for $(E, c, \mathcal{S}, (1+\delta)n, a)$ for some $\delta > 0$, then $C(x) \leq \frac{1}{\delta} C(x^*)$.

A Proof of Theorem 4.5

Proof of Theorem 4.5. We assume that $\alpha(\mathcal{C})$ is finite and omit the straightforward modifications necessary for classes with infinite α-value.

First suppose that for each scalar $\beta > 0$, there is some cost function $c \in \mathcal{C}$ satisfying $c(0) = \beta$. For any $\epsilon > 0$, choose a nonzero cost function $c \in \mathcal{C}$, a positive number $n > 0$ with $c(n) > 0$, and a scalar $\lambda \in [0,1]$ satisfying $c^*(\lambda n) = c(n)$ so that $[\lambda\mu + (1-\lambda)]^{-1} \geq \alpha(\mathcal{C}) - \epsilon/2$, where $\mu = c(\lambda n)/c(n)$. To mimic Example 2.4, we would like to employ a constant cost function everywhere equal to $c(n)$; however, this function need not lie in $\mathcal{C}$. Instead, we will choose $\tilde{c} \in \mathcal{C}$ satisfying $\tilde{c}(0) = c(n)$, and will use several "copies" of $\tilde{c}$ to "simulate" the function everywhere equal to $c(n)$. Toward this end, let m be so large that $\tilde{c}(\frac{(1-\lambda)n}{m-1}) \leq c(n) + \delta$ where δ is a sufficiently small positive number (depending on ϵ). Define a NCG Γ with $E = \{1, 2, \ldots, m\}$, one player type with singleton strategies $\mathcal{S}_1 = \{\{1\}, \{2\}, \ldots, \{m\}\}$ and $n_1 = n$, cost functions $c_1 = c$, $c_j = \tilde{c}$ for $j \in \{2, 3, \ldots, m\}$, and unit rates of consumption $a_{S,e} = 1$ for all S and e. The action distribution x with $x_{\{1\}} = n_1$ and $x_{\{j\}} = 0$ for $j > 1$ is an equilibrium for Γ with social cost $c(n)n$. The action distribution x^* with $x^*_{\{1\}} = \lambda n$ and $x_{\{j\}} = (1-\lambda)n/(m-1)$ for $j > 1$ has social cost at most $c(n)n[\lambda\mu + (1-\lambda) + \frac{1-\lambda}{c(n)}\delta]$. Provided δ is sufficiently small, $\rho(\Gamma) \geq \alpha(\mathcal{C}) - \epsilon$. Since $\epsilon > 0$ was arbitrary, the theorem holds for the class $\mathcal{C}$.

We now consider the general case. Suppose $\mathcal{C}$ is an inhomogeneous standard class of cost functions, and let $\widehat{\mathcal{C}}$ denote the closure of $\mathcal{C}$ under multiplication by positive scalars—that is, $\widehat{\mathcal{C}} = \{\nu c : c \in \mathcal{C}, \nu > 0\}$. Let $\widehat{\mathcal{G}}$ denote the set of NCGs with cost functions in $\widehat{\mathcal{C}}$. Since $\mathcal{C}$ is inhomogeneous, for any scalar $\beta > 0$ there is a function $\hat{c}$ in $\widehat{\mathcal{C}}$ with $\hat{c}(0) = \beta$. By the previous paragraph, $\sup_{\widehat{\Gamma} \in \widehat{\mathcal{G}}} \rho(\widehat{\Gamma}) \geq \alpha(\widehat{\mathcal{C}})$. Inspection of Definition 4.3 shows that $\alpha(\widehat{\mathcal{C}}) = \alpha(\mathcal{C})$. We can conclude the proof by showing that for any NCG $\widehat{\Gamma} \in \widehat{\mathcal{G}}$ and any $\epsilon > 0$, there is a NCG $\Gamma \in \mathcal{G}$ satisfying $\rho(\Gamma) \geq \rho(\widehat{\Gamma}) - \epsilon$.

Fix $\widehat{\Gamma} = (\widehat{E}, \hat{c}, \widehat{\mathcal{S}}, \hat{n}, \hat{a})$ and $\epsilon > 0$. For each element $\hat{e}$ of $\widehat{E}$ write $\hat{c}_{\hat{e}} = \nu_{\hat{e}} c_{\hat{e}}$ for $\nu_{\hat{e}} > 0$ and $c_{\hat{e}} \in \mathcal{C}$. The ratio ρ is a continuous function of each scalar $\nu_{\hat{e}}$, so we may replace each $\nu_{\hat{e}}$ by a sufficiently close positive rational number $\eta_{\hat{e}}$ to obtain a new NCG with ρ-value at least $\rho(\widehat{\Gamma}) - \epsilon$. Clearing denominators, we may assume that each scalar $\eta_{\hat{e}}$ is a positive integer (multiplying all cost functions of a NCG by a common positive number does not affect its ρ-value). The rest of the proof consists of observing that integral multiples of cost functions can be "simulated" with multiple elements possessing the original cost function. More precisely, define E by replacing each element $\widehat{e}$ of $\widehat{E}$ by $\eta_{\widehat{e}}$ new elements, each endowed with cost function $c_{\widehat{e}}$. Strategies $\mathcal{S}$ and rates of consumption a are then obtained in the obvious way from $\widehat{\mathcal{S}}$ and $\widehat{a}$ (with each element $\widehat{e}$ of $\widehat{E}$ in a strategy $\widehat{S}$ replaced by the $\eta_{\widehat{e}}$ corresponding elements of E, each with rate of consumption $\widehat{a}_{\widehat{S},\widehat{e}}$). It is straightforward to check that the natural bijective correspondence between action distributions of $(\widehat{E}, \eta c, \widehat{\mathcal{S}}, \widehat{n}, \widehat{a})$ and of $\Gamma = (E, c, \mathcal{S}, \widehat{n}, a)$ preserves both equilibria and social cost; therefore, $\Gamma \in \mathcal{G}$ with $\rho(\Gamma) \geq \rho(\widehat{\Gamma}) - \epsilon$, and the proof is complete. ■

References

Barro, R. J. and Romer, P. M. (1987). Ski-lift pricing, with applications to labor and other markets. *American Economic Review*, **77**(5), 875–890.

Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press.

Blonski, M. (1999). Anonymous games with binary actions. *Games and Economic Behavior*, **28**, 171–180.

Dafermos, S. C. and Sparrow, F. T. (1969). The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, **73B**(2), 91–118.

Friedman, E. J. (1996). Dynamics and rationality in ordered externality games. *Games and Economic Behavior*, **16**, 65–76.

Koutsoupias, E. and Papadimitriou, C. H. (1999). Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413.

Mas-Colell, A. (1984). On a theorem of Schmeidler. *Journal of Mathematical Economics*, **13**, 201–206.

Milchtaich, I. (1996). Congestion models of competition. *American Naturalist*, **147**(5), 760–783.

Milchtaich, I. (2000). Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research*, **25**(3), 349–364.

Milchtaich, I. (2001). Social optimality and cooperation in large congestion games. Manuscript.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, **14**, 124–143.

Papadimitriou, C. H. (2001). Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749–753.

Peressini, A. L., Sullivan, F. E., and Uhl, J. J. (1988). *The Mathematics of Nonlinear Programming*. Springer-Verlag.

Quint, T. and Shubik, M. (1994). A model of migration. Working paper, Cowles Foundation, Yale University.

Rapoport, A. and Chammah, A. (1965). *Prisoner's Dilemma*. University of Michigan Press.

Rath, K. P. (1992). A direct proof of the existence of pure strategy equilibria in games with a continuum of players. *Economic Theory*, **2**, 427–433.

Rosenthal, R. W. (1973a). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, **2**, 65–67.

Rosenthal, R. W. (1973b). The network equilibrium problem in integers. *Networks*, **3**, 53–59.

Roughgarden, T. (2002a). The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 428–437. Full version to appear in *Journal of Computer and System Sciences*.

Roughgarden, T. (2002b). *Selfish Routing*. Ph.D. thesis, Cornell University.

Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, **49**(2), 236–259.

Schmeidler, D. (1973). Equilibrium points of nonatomic games. *Journal of Statistical Physics*, **7**(4), 295–300.

Voorneveld, M., Borm, P., van Megan, F., Tijs, S., and Facchini, G. (1999). Congestion games and potentials reconsidered. *International Game Theory Review*, **1**, 283–299.

Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378.