

Worst-Case Efficiency Analysis of Queueing Disciplines^{*}

Damon Mosk-Aoyama^{**} and Tim Roughgarden^{1***}

Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305

1 Introduction

Consider n users vying for shares of a divisible good. Every user i wants as much of the good as possible but has diminishing returns, meaning that its *utility* $U_i(x_i)$ for $x_i \geq 0$ units of the good is a nonnegative, nondecreasing, continuously differentiable concave function of x_i . The good can be produced in any amount, but producing $X = \sum_{i=1}^n x_i$ units of it incurs a cost $C(X)$ for a given nondecreasing and convex function C that satisfies $C(0) = 0$. Cost might represent monetary cost, but other interesting interpretations are also possible. For example, x_i could represent the amount of traffic (measured in packets, say) that user i injects into a queue in a given time window, and $C(X)$ could denote aggregate delay ($X \cdot c(X)$, where $c(X)$ is the average per-unit delay). An altruistic designer who knows the utility functions of the users and who can dictate the allocation $x = (x_1, \dots, x_n)$ can easily choose the allocation that maximizes the *welfare* $W(x) = \sum_{i=1}^n U_i(x_i) - C(X)$, where $X = \sum_{i=1}^n x_i$, since this is a simple convex optimization problem.

But what if users are autonomous and choose the quantities x_i to maximize their own objectives? The most natural way to proceed is *equilibrium analysis*, where we model each user as maximizing its own payoff function and consider equilibrium allocations — those from which no user can unilaterally change its quantity to increase its payoff. We can then study the welfare achieved by autonomous and self-optimizing users via the *price of anarchy (POA)* — the worst (i.e., smallest) ratio between the welfare of an equilibrium (the outcome of selfish behavior) and the maximum-possible welfare (the ideal for an altruistic designer). The POA is a standard measure of inefficiency in game-theoretic systems, with a value near 1 indicating that selfish behavior is essentially benign.

Defining user payoffs requires a fundamental modeling decision: how does the joint cost $C(X)$ of producing an allocation translate to negative incentives for

^{*} The results in Section 3 of this work also appear in Chapter 5 of the first author's PhD thesis [4].

^{**} This work was done while the author was a PhD student at Stanford University, and supported in part by the ONR Young Investigator Award of the second author. Email: damonma@cs.stanford.edu.

^{***} Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship. Email: tim@cs.stanford.edu.

users? This choice is formalized by a *cost-sharing method* $\xi : \mathbf{R}_+^n \rightarrow \mathbf{R}_+^n$, which distributes the joint cost to the users: $\sum_i \xi_i(x) = C(X)$ for every allocation x with $X = \sum_i x_i$. For example, a natural cost-sharing method is *average-cost pricing*, defined by $\xi_i^{\text{FIFO}}(x) = \frac{x_i}{X} \cdot C(X)$; we also call this the “FIFO method.” In the queue example above, average-cost pricing naturally arises from the FIFO (first-in, first-out) queue service discipline with random packet arrivals. The other cost-sharing method that has been extensively studied in the present context is *serial cost-sharing*, which we define in Section 2 and also call the “Fair Share method,” after Shenker [11]. Given a cost-sharing method ξ , we define the payoff of user i in allocation x as $P_i(x) = U_i(x_i) - \xi_i(x)$; equilibria and the POA are then defined as outlined above. (Thus, we assume utility functions are in the same units as the cost function.)

From a design perspective, an obvious question is: *which cost-sharing method yields the best welfare guarantee (i.e., POA closest to 1)?* This question, as stated, is not well defined: the best cost-sharing method depends on the players’ utilities, and it is not reasonable to assume that these are known a priori to the designer. We therefore study the *worst-case POA* of cost-sharing methods, with the “worst” quantifier ranging over all possible utility function profiles U_1, \dots, U_n for a fixed number of users and a fixed cost function. Our research agenda is twofold: (1) For fundamental cost-sharing methods, precisely determine the worst-case POA in as many settings as possible; and (2) Identify the optimal cost-sharing method for a given environment — the one with the maximum-possible worst-case POA.

Our Results. Solving problems (1) and (2) in their full generality appears intractable, and our goal here is to provide precise answers for important special cases. Our first main result is for quadratic cost functions (of the form $C(X) = aX^2 + bX$). For a class of cost-sharing methods that strictly generalizes the FIFO and Fair Share methods, we give an exact formula for the worst-case POA of every method in the class for every number n of users. We give a single analysis that applies to all methods in the class. Our analysis identifies restricted linear structure in the equilibrium conditions for such methods and uses it to identify worst-case utility profiles. The precision of our analysis permits identification of the optimal method in this class for every number n of users. For example, our analysis shows that, in the limit as $n \rightarrow \infty$, the Fair Share method has the optimal worst-case POA among methods in the class.

General cost functions produce nonlinear equilibrium conditions and are much more difficult to analyze. For the Fair Share method and the case of $n = 2$ users, however, we show how to determine the worst-case POA with respect to general cost functions. This result is based on a novel and unexpected “reduction” to nonatomic selfish routing games.

Related Work. The work closest to ours is Moulin [5], who studies our exact model. Moulin [5] proved our first result for quadratic cost functions in the important special cases of the FIFO and Fair Share methods, using different (somewhat ad hoc) computations for each case. Here, we give a single analysis

generalizing these two results of his, which also applies to a broader class of cost-sharing methods. In our opinion, it is surprising that the seemingly very different FIFO and Fair Share methods (see Section 2) can be simultaneously analyzed with a common proof. Our result for general cost functions and two players generalizes a different result in Moulin [5], who gave tight bounds in the two-player case for monomial cost functions (for both the FIFO and Fair Share methods). The connection to selfish routing games is new to this paper, and it allows us to analyze general (non-monomial) cost functions. Moulin [5] also gave a number of results for the incremental cost-sharing method, which generally charges users more than the total production cost (i.e., is not “budget-balanced”) and therefore falls outside our purview. In subsequent work, Moulin [6] used budget-balanced cost-sharing methods with negative cost shares — subsidies, which are not permitted in this paper — to obtain much stronger positive results. For example, for every quadratic cost function and number of users, there is such a cost-sharing method that only induces games with POA equal to 1 [6].

A number of different but related models have been studied before. Closest to the present work is Johari and Tsitsiklis [2, 3], who studied a variant of our model with inelastic supply — i.e., there is a fixed amount of the divisible good but no production cost — and identified the allocation mechanism with the best worst-case POA among those in a broad class. We mention also Shenker [11], who studied the Fair Share method in a queueing context but without any quantitative efficiency guarantees; Moulin and Shenker [7], who compared the FIFO and Fair Share methods from an axiomatic perspective; and Christodoulou et al. [1], who were the first to study (in a different model) how to design protocols to optimize the worst-case POA.

2 Fundamental Cost-Sharing Methods

The FIFO method was defined in Section 1. The Fair Share method is an alternative designed to insulate users that request smaller quantities from the large requests. For example, with two players and quantities $x_1 \leq x_2$, the method assigns a cost share of $C(2x_1)/2$ to the first player (its fair share, if we pretend that the second player shares its size), and the balance $C(x_1 + x_2) - C(2x_1)/2$ to the second player. In general, all users split the cost that would ensue if all users were the same size as the smallest one; and the remaining cost is recursively allocated to the $n - 1$ largest users.

Precisely, using $[n]$ to mean $\{1, 2, \dots, n\}$: for a vector $x \in \mathbf{R}^n$, a permutation $\pi : [n] \rightarrow [n]$ is an *ordering* of x if the vector $z \in \mathbf{R}^n$ such that $z_{\pi(i)} = x_i$ satisfies $z_1 \leq z_2 \leq \dots \leq z_n$. The vector z is the *ordered version* of x . There are multiple orderings of a vector x when it has some equal components, but all the orderings give rise to the same ordered version z .

Definition 1 (The Fair Share Method [7, 11]). *For any cost function C , number of users n , and vector $x \in \mathbf{R}_+^n$, let π be an ordering of x with ordered version z . For $k \in [n]$, let $s_k = \sum_{\ell=1}^{k-1} z_\ell + (n - k + 1)z_k$. Then the cost share of*

user $i \in [n]$ is

$$\xi_i^{FS}(x) = \frac{C(s_{\pi(i)})}{n - \pi(i) + 1} - \sum_{k=1}^{\pi(i)-1} \frac{C(s_k)}{(n - k + 1)(n - k)}.$$

A simple way to interpolate between the FIFO and Fair Share methods is via the following θ -combinations for $\theta \in [0, 1]$:

$$\xi_i(x) = \theta \xi_i^{FS}(x) + (1 - \theta) \xi_i^{FIFO}(x). \quad (1)$$

The FIFO and Fair Share methods correspond to the values $\theta = 0$ and $\theta = 1$, respectively. A θ -combination can be implemented in a system with one FIFO queue and one Fair Share queue. Each arriving packet is placed into the Fair Share queue with probability θ (and otherwise the FIFO queue). Departing packets are chosen from a queue with these same probabilities. We emphasize that while θ -combinations are defined as a linear combination of the FIFO and Fair Share methods, the equilibria and POA with respect to such methods are *not* linear in θ — indeed, even for a quadratic cost function, the worst-case POA is a fairly complex function of θ (Theorem 1).

3 Quadratic Cost Functions

This section considers quadratic cost functions. This assumption is restrictive, but we will be rewarded with an exact characterization of the worst-case price of anarchy for every θ -combination and number n of users. For simplicity, we assume that $C(X) = X^2$ throughout this section; scaling by a constant changes nothing, and adding a linear term (with a nonnegative coefficient) only improves the POA.

3.1 Equilibrium Properties

We now state some basic properties of equilibria with respect to a θ -combination and the cost function $C(X) = X^2$. The proofs of these preliminary results are not trivial, but they are not the main point of this paper and we refer the interested reader to [4, Chapter 5] for the technical details.

For a given θ -combination, with a quadratic cost function, there is a linear relationship between an allocation vector x and the corresponding marginal costs $\xi'_i(x_i; x_{-i})$. (Here the derivative is w.r.t. x_i and x_{-i} denotes the other users' quantities — the vector x with the i th component removed.) For example, Figure 1 demonstrates this linear relationship in the special case of the FIFO and Fair Share methods, when $n = 4$. The next proposition formalizes the relationship for all of the cost-sharing methods that we study.

$$B^0 = \begin{bmatrix} 2 & 1 & 1 & 1 \\ 1 & 2 & 1 & 1 \\ 1 & 1 & 2 & 1 \\ 1 & 1 & 1 & 2 \end{bmatrix} \quad B^1 = \begin{bmatrix} 8 & 0 & 0 & 0 \\ 2 & 6 & 0 & 0 \\ 2 & 2 & 4 & 0 \\ 2 & 2 & 2 & 2 \end{bmatrix}$$

Fig. 1. The matrix B defined in Proposition 1 for the FIFO method (B^0) and Fair Share method (B^1), when $n = 4$. Assuming that the users have been sorted in nondecreasing order of x_i , the columns correspond to the quantities x_i , and the rows to the induced marginal costs $\xi'_i(x_i; x_{-i})$.

Proposition 1. *For the cost function $C(y) = y^2$, any number of users n , any user $i \in [n]$, and any $\theta \in [0, 1]$, let ξ_i denote the cost share of user i under the θ -combination. Define an $n \times n$ matrix B by*

$$B_{k\ell} = \begin{cases} 2(1 + \theta(n - k)), & \text{if } k = \ell; \\ 1 + \theta, & \text{if } k > \ell; \\ 1 - \theta, & \text{if } k < \ell \end{cases}$$

for any $k, \ell \in [n]$. For any vector $x \in \mathbf{R}_+^n$, let π be any ordering of x , and let z be the ordered version of x .

- (a) The vector $p \in \mathbf{R}^n$ with $p_{\pi(i)} = \xi'_i(x_i; x_{-i})$ for all $i \in [n]$ is given by $p = Bz$.
- (b) If $z_{k_1} < z_{k_2}$ then $p_{k_1} < p_{k_2}$; if $z_{k_1} \leq z_{k_2}$ then $p_{k_1} \leq p_{k_2}$.

Part (a) of Proposition 1 asserts that the matrix B correctly maps allocations to marginal costs. Part (b) asserts that marginal costs must be increasing in the quantities x_i . Proposition 1 will be useful in our POA analysis and also enables us to establish existence and uniqueness of equilibria.

Proposition 2. *For every θ -combination, quadratic cost function, and utility function profile, the induced game has a unique equilibrium allocation vector.*

Proposition 2 also holds for much more general convex cost functions. It can be proved by modifying Rosen's existence and uniqueness theorems for convex games [8]. (Modifications are needed because the Fair Share method is not continuously differentiable at allocation vectors with two equal components.)

3.2 Tight Bounds on the POA

We now show how to determine the worst-case price of anarchy of every θ -combination under a quadratic cost function, over all utility function profiles. We first note that linear functions — of the form $U_i(x_i) = a_i x_i$ for $a_i \geq 0$ — induce games with as large a POA as any other type of (nonnegative and concave) utility function.

Lemma 1 (Linearization Lemma [2, 5]). *For every θ -combination, number n of users, and convex cost function, the worst-case POA (over all utility function profiles) is determined by linear utility function profiles.*

The proof of Lemma 1 simply shows that linearizing utility functions at the equilibrium point only worsens the POA, and then shifting the resulting nonnegative affine functions to be linear again only worsens the POA.

For the rest of this section, we assume that all utility functions are linear with $0 < a_1 \leq a_2 \leq \dots \leq a_n$. For such a profile, an optimal solution allocates only to the n th user. A simple calculation shows that the optimal amount to give this user is $a_n/2$, leading to a welfare of $a_n(a_n/2) - (a_n/2)^2 = a_n^2/4$.

The next lemma studies the welfare of an equilibrium allocation x^* and is central to our analysis. It states that the requested quantities in x^* are in the same order as the a_i values, determines a remarkable formula for the welfare of the system under x^* , and develops a constraint that relates the x_i^* values to a_n .

Lemma 2. *For the cost function $C(y) = y^2$, any number of users n , any $\theta \in [0, 1]$, and any $a \in \mathbf{R}_+^n$ such that $0 < a_1 \leq a_2 \leq \dots \leq a_n$, let x^* be the equilibrium allocation under the θ -combination.*

- (a) *The requested quantities in x^* are in the order $x_1^* \leq x_2^* \leq \dots \leq x_n^*$.*
- (b) *The welfare of the system under x^* is*

$$W(x^*) = \sum_{i=1}^n (2\theta(n-i) + 1) (x_i^*)^2. \quad (2)$$

- (c) *The components of x^* satisfy the equation*

$$(1 + \theta) \sum_{i=1}^{n-1} x_i^* + 2x_n^* = a_n. \quad (3)$$

Proof. Simple computations show that, for every i , the function $\xi_i(y; x_{-i})$ is convex and differentiable in y for every fixed x_{-i} . It follows that an allocation vector x is an equilibrium, with each user i choosing an optimal quantity given x_{-i} , if and only if

$$\begin{aligned} x_i^* > 0 &\Rightarrow a_i = \xi_i'(x_i^*; x_{-i}^*); \\ x_i^* = 0 &\Rightarrow a_i \leq \xi_i'(0; x_{-i}^*). \end{aligned} \quad (4)$$

To prove (a), let x^* be an equilibrium and suppose for contradiction that there are two users i_1 and i_2 such that $i_1 < i_2$ and $x_{i_1}^* > x_{i_2}^*$. Because $x_{i_2}^* \geq 0$, we have $x_{i_1}^* > 0$, and so the equilibrium conditions in (4) imply that $a_{i_1} = \xi_{i_1}'(x_{i_1}^*; x_{-i_1}^*)$ and $a_{i_2} \leq \xi_{i_2}'(x_{i_2}^*; x_{-i_2}^*)$. Since $a_{i_1} \leq a_{i_2}$, $\xi_{i_1}'(x_{i_1}^*; x_{-i_1}^*) \leq \xi_{i_2}'(x_{i_2}^*; x_{-i_2}^*)$. On the other hand, since $x_{i_1}^* > x_{i_2}^*$, Proposition 1 implies that $\xi_{i_1}'(x_{i_1}^*; x_{-i_1}^*) > \xi_{i_2}'(x_{i_2}^*; x_{-i_2}^*)$, contradicting this inequality.

Given that $x_1^* \leq x_2^* \leq \dots \leq x_n^*$, we can apply Proposition 1 with the ordering π of x^* being the identity permutation to rewrite the equilibrium conditions in (4) as

$$\begin{aligned} x_i^* > 0 &\Rightarrow a_i = \sum_{j=1}^n B_{ij}x_j^*; \\ x_i^* = 0 &\Rightarrow a_i \leq \sum_{j=1}^n B_{ij}x_j^*, \end{aligned} \tag{5}$$

where B is the $n \times n$ matrix defined in Proposition 1. By definition, the welfare of x^* is

$$W(x^*) = \sum_{i=1}^n a_i x_i^* - \left(\sum_{i=1}^n x_i^* \right)^2. \tag{6}$$

The equilibrium conditions in (5) imply that total utility at equilibrium can be written as a quadratic form:

$$\sum_{i=1}^n a_i x_i^* = \sum_{i=1}^n x_i^* \sum_{j=1}^n B_{ij} x_j^* = (x^*)^T B x^* = (x^*)^T \left(\frac{1}{2} (B + B^T) \right) x^*.$$

Since a quadratic cost function can be similarly expressed as

$$C(x) = \left(\sum_i x_i \right)^2 = x^T E x,$$

where E is the all-ones $n \times n$ matrix, *equilibrium welfare can be expressed as a quadratic form*:

$$\begin{aligned} W(x^*) &= (x^*)^T \left(\frac{1}{2} (B + B^T) \right) x^* - (x^*)^T E x^* \\ &= (x^*)^T \left(\frac{1}{2} (B + B^T) - E \right) x^*. \end{aligned} \tag{7}$$

Let D denote the symmetric matrix $\frac{1}{2}(B + B^T) - E$. By the definition of B , the diagonal entries of D are $D_{ii} = B_{ii} - 1 = 2\theta(n - i) + 1$ for all $i \in [n]$. For any $i, j \in [n]$ such that $i \neq j$, we have $D_{ij} = D_{ji} = (1/2)(1 + \theta + 1 - \theta) - 1 = 0$. Thus D is a diagonal matrix, and the equation involving the quadratic form $(x^*)^T D x^*$ in (7) simplifies to

$$W(x^*) = \sum_{i=1}^n D_{ii} (x_i^*)^2,$$

which yields the expression in (2) upon substitution of the D_{ii} values.

Finally, since $a_n > 0$, the equilibrium condition in (5) implies that $x_n^* > 0$ with

$$a_n = \sum_{i=1}^n B_{ni} x_i^*.$$

By substituting the B_{ni} values, we obtain the equation in (3). \square

Scaling a vector of coefficients a by λ increases both the optimal and equilibrium welfares by a factor of λ^2 (for the latter, this follows from the linear equilibrium conditions and Lemma 2(b)). Since the POA is the ratio of these, and since the optimal welfare depends only on a_n , we can restrict our search for the worst-case utility function profile to the set $\mathcal{A} = \{a \in \mathbf{R}_+^n \mid 0 < a_1 \leq a_2 \leq \dots \leq a_n = 1\}$ and focus on minimizing the equilibrium welfare W_a over $a \in \mathcal{A}$. Our second key lemma computes this minimum precisely.

Lemma 3. *Fix the cost function $C(y) = y^2$, any number of users n , and any $\theta \in [0, 1]$. Then $\inf_{a \in \mathcal{A}} W_a = 1/4\Gamma_\theta(n)$, where W_a is the welfare of the (unique) equilibrium for the utility profile a , and*

$$\Gamma_\theta(n) = 1 + \frac{(1 + \theta)^2}{4} \sum_{i=1}^{n-1} \frac{1}{2\theta i + 1}. \quad (8)$$

Proof. By Lemma 2, a lower bound on the minimum-possible equilibrium welfare is provided by the value of the convex program

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^n (2\theta(n-i) + 1)x_i^2 \\ & \text{subject to} && (1 + \theta) \sum_{i=1}^{n-1} x_i + 2x_n = 1. \end{aligned} \quad (9)$$

We introduce a Lagrange multiplier λ for the constraint $1 - (1 + \theta) \sum_{i=1}^{n-1} x_i - 2x_n = 0$. Then the Karush-Kuhn-Tucker (KKT) optimality conditions for the program in (9) are

$$\begin{aligned} 2(2\theta(n-i) + 1)x_i - \lambda(1 + \theta) &= 0, \quad \forall i \in [n-1]; \\ 2x_n - 2\lambda &= 0. \end{aligned}$$

Solving the KKT conditions for the x_i values yields

$$\begin{aligned} x_i &= \left(\frac{1 + \theta}{2\theta(n-i) + 1} \right) \frac{\lambda}{2}, \quad \forall i \in [n-1]; \\ x_n &= \lambda. \end{aligned}$$

Substituting these values into the equality constraint in (9), we obtain

$$\begin{aligned} 1 &= \lambda \left(\frac{(1 + \theta)^2}{2} \sum_{i=1}^{n-1} \frac{1}{2\theta(n-i) + 1} + 2 \right) \\ &= 2\lambda \left(\frac{(1 + \theta)^2}{4} \sum_{i=1}^{n-1} \frac{1}{2\theta i + 1} + 1 \right) \\ &= 2\lambda\Gamma_\theta(n), \end{aligned}$$

and thus $\lambda = 1/2\Gamma_\theta(n)$.

The value of the objective function in (9) for this vector x is

$$\begin{aligned} \sum_{i=1}^{n-1} (2\theta(n-i) + 1) \left(\frac{1+\theta}{2\theta(n-i) + 1} \right)^2 \left(\frac{\lambda}{2} \right)^2 + \lambda^2 &= \lambda^2 \left(\frac{(1+\theta)^2}{4} \sum_{i=1}^{n-1} \frac{1}{2\theta(n-i) + 1} + 1 \right) \\ &= \lambda^2 \Gamma_\theta(n) \\ &= \frac{1}{4\Gamma_\theta(n)}. \end{aligned}$$

This quantity lower bounds the minimum-possible equilibrium welfare (for vectors $a \in \mathcal{A}$).

To obtain a matching upper bound, consider the vector $x \in \mathbf{R}_+^n$ obtained by solving the KKT conditions and imposing the equality constraint in (9). The components of x are

$$\begin{aligned} x_i &= \left(\frac{1+\theta}{2\theta(n-i) + 1} \right) \frac{1}{4\Gamma_\theta(n)}, \quad \forall i \in [n-1]; \\ x_n &= \frac{1}{2\Gamma_\theta(n)}, \end{aligned}$$

and so $0 \leq x_1 \leq x_2 \leq \dots \leq x_n$. Define a vector a as $a = Bx$, where B is the matrix defined in the statement of Proposition 1 for the θ -combination. Proposition 1 implies that $a_1 \leq a_2 \leq \dots \leq a_n$. Moreover, because x satisfies the equality constraint in (9), we have

$$a_n = \sum_{i=1}^{n-1} B_{ni}x_i + B_{nn}x_n = (1+\theta) \sum_{i=1}^{n-1} x_i + 2x_n = 1.$$

Finally, the vector x satisfies the equilibrium conditions in (5), so it is the equilibrium of the game with utility functions defined by a . Since the welfare of x is $W_a = 1/4\Gamma_\theta(n)$, the proof is complete. \square

Recalling that the optimal welfare is $1/4$ for every vector $a \in \mathcal{A}$, we have our main result for quadratic cost functions.

Theorem 1. *For the cost function $C(y) = y^2$, any number of users n , and any $\theta \in [0, 1]$, the price of anarchy of the θ -combination is $1/\Gamma_\theta(n)$.*

The formula in Theorem 1 for the two special cases $\theta = 0$ and $\theta = 1$ was established in Moulin [5] using two different proofs. Obviously, the formula in (8) can be used to identify the optimal θ -combination for every number n of users. When $n = 2$, the FIFO method is the best θ -combination (with worst-case POA $4/5$) while the Fair Share method is the worst (with worst-case POA $3/4$). When $n = 3$, the Fair Share method remains the worst (with worst-case POA $15/23$). The FIFO method is slightly better, with worst-case POA $2/3$. Taking $\theta \approx .262648$ yields a superior θ -combination, with worst-case POA $\approx .687$. For

n	POA of FIFO	POA of Fair Share	Optimal θ	Optimal POA
2	.8	.75	0	.8
3	$\approx .667$	$\approx .652$	$\approx .262648$	$\approx .687$
4	$\approx .571$	$\approx .595$	$\approx .375361$	$\approx .623$
5	.5	$\approx .559$	$\approx .442921$	$\approx .581$
10	$\approx .308$	$\approx .469$	$\approx .588111$	$\approx .481$
20	$\approx .174$	$\approx .403$	$\approx .677465$	$\approx .410$
40	$\approx .093$	$\approx .354$	$\approx .737$	$\approx .358$

Table 1. Illustration of Theorem 1. Comparison, for different n , of the exact worst-case POA of the FIFO method, the Fair Share method, and the optimal θ -combination.

all $n \geq 4$, the Fair Share method outperforms the FIFO method but other θ -combinations are still better (see Table 1). In the limit as $n \rightarrow \infty$, for every $\theta \in (0, 1]$, $\Gamma_\theta(n)$ scales as $(1 + \theta)^2 \ln n / (8\theta)$. (For $\theta = 0$, the FIFO method scales as $4/(n + 3)$.) Since $\theta/(1 + \theta)^2$ is increasing in the interval $\theta \in (0, 1]$, the Fair Share method has the best asymptotic worst-case POA, which scales as $2/(\ln n)$ for large n .

4 General Cost Functions

This section considers general (non-quadratic) cost functions. Analyzing the case of many players appears intractable, so we settle for a solution to two-player games induced by the Fair Share method. We begin with a simple but useful lemma, which holds even with many users.

Lemma 4. *If all users have linear utility functions and the cost function is strictly convex, then the total quantity allocated in the Fair Share equilibrium equals that in the optimal allocation.*

Proof. (Sketch.) Suppose user n has the largest utility function coefficient a_n . One optimal solution allocates only to user n , and the optimal amount to allocate is the unique point X at which $a_n = C'(X)$.

Consider an equilibrium under Fair Share. Analogous to condition (5) in Lemma 2 and the bottom row of the matrix B^1 in Figure 1, at equilibrium we must have $a_n = \xi'_n(x_n^*; x_{-n}^*) = C'(X)$. Thus the quantity allocated at x^* (across all users) equals that in the optimal solution (to user n only). \square

Our approach is to show an explicit connection, interesting in its own right, between games with two users with linear utility functions and nonatomic selfish routing games (e.g. [10]). Recall *Pigou's example*, a basic selfish routing network: X units of traffic, comprising a large number of infinitesimal autonomous users, choose between two parallel links connecting one vertex to another. One link has some per-unit cost function $c(x_1)$, and the other has constant per-unit cost $c(X)$. The first link is a dominant strategy, so in the only equilibrium all traffic takes it and the aggregate cost is $X \cdot c(X)$. An optimal outcome, by definition, splits the

traffic x_1 and $x_2 = X - x_1$ between the two links to minimize $x_1 \cdot c(x_1) + x_2 \cdot c(X)$. Much is known about the ratio between the equilibrium and optimal costs in Pigou's example (and much more general selfish routing networks), as a function of c . This ratio is called the *Pigou bound for c* and is denoted $\alpha(c)$. For example, the Pigou bound for all affine functions is at most $4/3$, with equality achieved when $c(x) = ax$ for some $a > 0$; and for per-unit cost functions c that are polynomials with nonnegative coefficients and degree at most p , the largest Pigou bound grows like $\approx p/\ln p$ [9].

The connection between Pigou's example and the queueing games studied in this paper is most vivid for the total user utility — so for our penultimate result, we ignore the cost term in the welfare objective function.

Theorem 2. *For every differentiable convex cost function C , the worst-case fraction of the optimal total utility achieved by the Fair Share equilibrium allocation with two players is exactly*

$$\frac{1}{2} \left(1 + \frac{1}{\alpha(C')} \right),$$

where $\alpha(C')$ is the Pigou bound for C' .

Proof. (Sketch.) We prove the theorem constructively, by exhibiting a worst-possible example for the Fair Share equilibrium allocation. Fix a choice of $X \geq 0$; we later optimize adversarially over X . Give the second user the utility function $U_2(x_2) = C'(X) \cdot x_2$. For any $x_1 \in [0, X/2]$, choosing the coefficient $a_1 = C'(2x_1) \leq C'(X)$ for the first user's utility function ensures that $x_1^* = x_1$ at the Fair Share equilibrium x^* . For a given choice of X and $x_1 \in [0, X/2]$, the total user utility obtained by Fair Share is then $x_1 \cdot C'(2x_1) + (X - x_1) \cdot C'(X)$. By a change of variable, this payoff is minimized at $y^*/2$, where y^* is the optimal amount of traffic to route on the non-constant link of Pigou's example when there are X units of traffic and the non-constant per-unit cost function c is C' . The resulting total user utility is

$$\begin{aligned} \frac{y^*}{2} C'(y^*) + \left(X - \frac{y^*}{2} \right) \cdot C'(X) &= \frac{1}{2} (y^* C'(y^*) + (X - y^*) \cdot C'(X)) + \frac{1}{2} (X \cdot C'(X)) \\ &\geq X \cdot C'(X) \left(\frac{1}{2} + \frac{1}{2\alpha(C')} \right), \end{aligned}$$

where the inequality follows from the definition of the Pigou bound $\alpha(C')$ for C' . The parameter X can be chosen so that the inequality is arbitrarily close to an equality. The total user utility obtained in the optimal solution is $X \cdot C'(X)$, and the Fair Share equilibrium allocation obtains only a $\frac{1}{2}(1 + 1/\alpha(C'))$ fraction of this. Reversing the steps in the argument above shows that no worse example is possible. \square

To extend Theorem 2 to a bound on the POA for the welfare objective, we need to re-introduce the cost terms (for both the optimal and equilibrium

allocations). This can be approached in a number of ways. Since we can assume utility functions are linear (Lemma 1), Lemma 4 shows that both allocations will suffer the same cost. A crude way to proceed is to define

$$\gamma(c) = \sup_{X \geq 0} \frac{\int_0^X c(x) dx}{X \cdot c(X)}; \quad (10)$$

for example, if $c(x) = x^d$, then $\gamma(c) = 1/(d + 1)$. Theorem 2 then yields the following corollary for the POA.

Corollary 1. *For every differentiable convex cost function C , the worst-case POA of Fair Share with two players is at least*

$$\frac{1}{1 - \gamma(C')} \cdot \left(\frac{1}{2} \left(1 + \frac{1}{\alpha(C')} \right) - \gamma(C') \right),$$

where $\alpha(C')$ is the Pigou bound for C' and $\gamma(C')$ is defined as in (10).

For example, for the marginal cost function $C'(x) = x^d$, plugging in the known upper bound on the Pigou value [9] together with $\gamma(C') = 1/(d + 1)$ immediately gives a lower bound of $1 - \frac{1}{2}(d + 1)^{-1/d}$, on the worst-case POA, recovering a result of Moulin [5]. Other natural types of cost functions can be treated in a similar way.

References

1. G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. In *Proceedings of the 31st Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3142 of *Lecture Notes in Computer Science*, pages 345–357, 2004.
2. R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operation Research*, 29(3):407–435, August 2004.
3. R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 2009. To appear.
4. D. Mosk-Aoyama. *Convergence to and Quality of Equilibria in Distributed Systems*. PhD thesis, Stanford University, 2008.
5. H. Moulin. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory*, 36(3):379–405, September 2008.
6. H. Moulin. An efficient and almost budget balanced cost sharing method. *Games and Economic Behavior*, 2009. In press.
7. H. Moulin and S. J. Shenker. Serial cost sharing. *Econometrica*, 60(5):1009–1037, September 1992.
8. J. B. Rosen. Existence and uniqueness of equilibrium points for concave n -person games. *Econometrica*, 33(3):520–534, July 1965.
9. T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
10. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
11. S. J. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3(6):819–831, December 1995.