

Mediated Equilibria in Load-Balancing Games*

Joshua R. Davis David Liben-Nowell Alexa Sharp Tom Wexler

November 6, 2011

Abstract: *Mediators* are third parties to whom the players in a game can delegate the task of choosing a strategy; a mediator forms a *mediated equilibrium* if delegating is a best response for all players. Mediated equilibria have more power to achieve outcomes with high social welfare than Nash or correlated equilibria, but less power than a fully centralized authority. Here we initiate the study of the power of mediation by introducing the mediation analogue of the price of stability—the ratio of the social cost of the best mediated equilibrium BME to that of the socially optimal outcome OPT. We focus on load-balancing games with social cost measured by weighted average latency. Even in this restricted class of games, BME can range from being as good as OPT to being no better than the best correlated equilibrium. In unweighted games BME achieves OPT; the weighted case is more subtle. Our main results are (1) a proof that the worst-case ratio BME/OPT is at least $(1 + \sqrt{2})/2 \approx 1.2071$ and at most 2 for linear-latency weighted load-balancing games, and that the lower bound is tight when there are two players; and (2) tight bounds on the worst-case BME/OPT for general-latency weighted load-balancing games. We give similarly detailed results for other natural social-cost functions. We also show a precise correspondence between the quality of mediated equilibria in a game G and the quality of Nash equilibria in the repeated game when G is played as the stage game. The proof of this correspondence is based on a close connection between mediated equilibria and the Folk Theorem from the economics literature.

Key words and phrases: mediated equilibria, load-balancing games, quality of equilibria

*A preliminary abbreviated version of this paper appears in the *Proceedings of the 5th Workshop on Internet and Network Economics (WINE'09)*, pp. 591–599, December 2009. All authors were supported in part by NSF grant CCF-0728779. This work was also supported by grants from Carleton College, Oberlin College, and Denison University. The work of the fourth author was performed in part at Denison University. Thanks to Bobby Kleinberg for suggesting this line of inquiry to us, to Ioannis Caragiannis for helpful discussions and for communicating his proof of Theorem 4.12, and to Lauren Feiler and the anonymous referees for very helpful suggestions.

1 Introduction

The recent interest in algorithmic game theory by computer scientists is in large part motivated by the recognition that the implicit assumptions of traditional algorithm design are ill-suited to many real-world settings. Algorithms are typically designed to transform fully specified inputs into solutions that can then be implemented by some centralized authority. But in many situations—the routing of traffic, the creation of large-scale networks, or the operation of networks, for example—no such centralized authority exists; solutions arise through the interactions of self-interested, independent agents. Thus researchers have begun to use game theory to model these competitive, decentralized situations.

One classic and influential example is the paper of Koutsoupias and Papadimitriou [24], who consider the effect of decentralizing a standard load-balancing problem. In the resulting game, each job is controlled by a distinct player who selects a machine to serve her job so as to minimize her own delay. The authors compare the social cost (expected maximum delay) of the Nash equilibria of this game to that of a centrally designed optimal solution. The maximum of these ratios is the *price of anarchy* of the game, quantifying the worst-case cost of decentralized behavior: how much more expensive is the uncoordinated equilibrium than the coordinated optimum?

Although we have cast the contrast between “centralized” and “decentralized” solutions starkly, one can imagine a continuum indexing the amount of power that a centralized authority has in implementing solutions to a given problem. At one extreme is a centralized authority with complete control: a benevolent dictator, who can implement a globally optimal solution. At the other extreme is utter impotence: with no centralized power, we would expect to reach a solution that is a Nash equilibrium, and potentially the worst of these equilibria. The price of anarchy effectively measures the difference between the two extremes of this continuum. Of course, just as it is unrealistic to assume that an omnipotent centralized authority always exists, it is also implausible to postulate a complete absence of centralized control; intermediate points in this continuum are also interesting.

For example, we can consider a weak centralized authority who can *propose* a solution simultaneously to all the players, but has no power to force the players to take any particular action in response to that proposal. The players would all be willing to follow such a proposal if and only if it were a Nash equilibrium—but the authority could propose the *best* Nash equilibrium (rather than the worst, which we assumed would arise in the complete absence of centralization). The ratio of the cost of the best Nash equilibrium to the global optimum is the *price of stability* of the game. In some settings, even this small amount of power is significant; the price of stability can be dramatically smaller than the price of anarchy.

We can grant additional power to the centralized authority by eliminating the requirement of broadcasting the entirety of a proposed solution to all players. This additional power leads to a broader notion of equilibrium, the *correlated equilibrium* [3]. Such a centralized authority, or *correlator*, may signal each player individually with a suggested action, chosen by the correlator according to some known joint probability distribution ψ . Outcomes that are not stable under the definition of a Nash equilibrium may be stable given the limited information that each player individually receives. A correlator is a genuinely more powerful designer: any Nash equilibrium is a correlated equilibrium, but often a correlator can induce outcomes that are substantially better than the best Nash equilibrium.

Mediators and mediated equilibria. This paper focuses on centralized authorities known as *mediators* [1, 27, 29–32, 35], which are endowed with even more power than correlators. A mediator is a centralized authority who offers to act on behalf of any player who chooses to use the mediator’s services, by choosing a strategy in that player’s stead. Each player has the option of *delegating* control to the mediator, who will play simultaneously on behalf of all players who have chosen this option. The strategies that the mediator chooses for the delegating players may be correlated; furthermore, the distribution from which the mediator draws these strategies depends on which players have opted to use mediation. These distributions are known to all players: that is, a mediator is defined by a distribution ψ_T over tuples of strategies for the players in T , for every subset T of players. The mediator $\{\psi_T : T \text{ is a subset of the players}\}$ forms a *mediated equilibrium* if all players prefer to delegate than to play on their own behalf. The fact that ψ_T depends on T is significant, because it endows the mediator with extra influence in discouraging players from leaving mediation. In particular, a player who decides to leave the mediator must consider that her departure from the mediator changes the set of delegating players, and thus may change the distribution of strategies for the remaining delegating players as well. In certain games and certain mediated equilibria, the mediator is able to persuade a player to stick with mediation by threatening to have the remaining players “punish” the defecting player were she to leave. However, it is important to note that mediators are still limited in their power: they have no way to enforce any particular action by any player.

It is not difficult to show that any correlated equilibrium can be represented as a mediated equilibrium, while there can be mediated equilibria that do not correspond to any correlated equilibrium. Thus mediators are in general more powerful than correlators. It is also worth noting that while a mediated equilibrium corresponds to a solution concept that is distinct from Nash equilibria, we can nevertheless describe mediation in terms of Nash equilibria. In particular, mediated equilibria in any game can be thought of as Nash equilibria in an extension of the original game, in which each player has one additional strategy available (namely, to delegate control to the mediator). The payoffs of the new outcomes in which some players choose to delegate depend both on the payoffs of the original game, as well as the particular mediator being used.

The present work: mediators in load-balancing games. In this paper, we begin the task of quantifying the powers and limitations of mediators. We consider the mediation analogue of the price of stability: how much less efficient than the globally optimal outcome is the best mediated equilibrium? (While one could ask questions analogous to the price of anarchy instead, the spirit here is that of a well-intentioned centralized authority who would aim for the best, not the worst, outcome within its power.) We initiate this study in the context of *load-balancing games*. Each player i controls a job that must be assigned to a machine for processing. Each machine j is characterized by a nonnegative and nondecreasing latency function $f_j(x)$, and each player incurs a *cost* of $f_j(\ell_j)$ for choosing machine j , where ℓ_j denotes the total load of jobs that choose machine j . We split load-balancing games into classes along two dimensions:

- *unweighted* versus *weighted*: in weighted load-balancing games, each job i has weight w_i , and experiences a cost on machine j of $f_j(\sum_{i' \text{ uses } j} w_{i'})$; in unweighted games all $w_i = 1$.
- *linear* versus *general* latencies: in linear games, $f_j(x) = a_j \cdot x$ for $a_j \geq 0$; for general latencies the function f_j can be an arbitrary nonnegative and nondecreasing function.

The social cost is measured by the weighted average latency experienced by the jobs; see Section 6 for results using other social-cost functions.

Load-balancing games form an appealing domain for this work, for two reasons. First, they form a simple yet rich class of games for this type of analysis—which includes cases in which mediators can achieve OPT and cases in which they cannot even do better than the best Nash equilibrium BNE. Second, the prices of anarchy and stability, and corresponding measures of correlated equilibria, are well understood for these games and many of their variants [5, 8–10, 14, 23–25, 34]. Most relevant for what follows are an upper bound of 2 on the price of stability in weighted linear games due to Caragiannis [7] (based on techniques of Fotakis–Kontogiannis–Spirakis [14]; see Theorem 4.12) and a tight upper bound of $4/3$ on the price of stability in unweighted linear games [8]. We extend this line of work to mediated equilibria with the following results:

- In the unweighted case, the best mediated equilibrium BME is optimal, regardless of the latency functions’ form. This result follows from a recent theorem of Monderer and Tennenholtz [27], which in fact holds for any symmetric game. See Section 3.
- In weighted linear-latency games with two players, we give tight bounds on the best solution a mediator can guarantee: at most a factor of $(1 + \sqrt{2})/2 \approx 1.2071$ worse than OPT but up to $4/3$ better than the best correlated equilibrium BCE. Thus mediators cannot always implement optimal solutions, but are strictly more powerful than correlators. See Section 4.
- In weighted nonlinear-latency games, mediated equilibria provide no worst-case improvement over correlated or even Nash equilibria. See Section 5.
- We also analyze mediation under two other social-cost functions that are considered in the literature: (i) the maximum latency of the jobs; and (ii) the average latency, unweighted by the jobs’ weights. Under (i), the price of stability is one, and thus the additional power of a mediator is uninteresting. Under (ii), we show qualitatively similar results to those described above for the weighted average social-cost function; we also show that the best mediated equilibrium can be unboundedly worse than the social optimum in linear-latency games as the number of players grows. See Section 6.

Our results for social costs measured by weighted average latency are summarized in Figure 1. Because mediated equilibria naturally fall between correlated equilibria and full centralized control on the continuum, we compare the best mediated equilibrium BME to the best correlated equilibrium BCE and the global optimum OPT in each case.

The notion of punishment from mediated equilibria is closely related to the Folk Theorem from the economics literature. (See, for example, the classical game theory text of Fudenberg and Tirole [17].) Roughly, the Folk Theorem states that in a repeated game with sufficiently patient players, any feasible, individually rational outcome s is an equilibrium. The typical proof is based on the idea of a “grim trigger”: if a player i ever chooses a strategy other than s_i , then in every future round of the game, every other player plays according to a punishing strategy that causes player i to incur very high costs. As long as s gives a higher-quality outcome to player i , then this threat of future heavy punishment is enough to cause each player to adhere to s in every round. Under very mild assumptions, we show a tight correspondence between the ratio BME/OPT in a game G and the price of stability (the ratio of the best Nash equilibria BNE to OPT) when the game G is the stage game of a repeated game.

	unweighted jobs	weighted jobs
linear latencies	BME = OPT [27] BCE $\leq 4/3 \cdot$ BME [tight] (Lem. 3.6 [8])	BME $\leq 2 \cdot$ OPT (Thm. 4.12 [7]) BCE $\leq 2 \cdot$ BME (Thm. 4.12 [7]) $n = 2$: BME $\leq \frac{1+\sqrt{2}}{2} \cdot$ OPT (Thm. 4.1) BCE $\leq 4/3 \cdot$ BME (Thm. 4.1) [both tight for $n = 2$]
general latencies	BME = OPT [27] BCE $\leq n \cdot$ BME [tight] (Lem. 3.6)	BME $\leq \Delta \cdot$ OPT [tight] (Thm. 5.1) BCE $\leq \Delta \cdot$ BME [tight] (Thm. 5.1)

Figure 1: Summary of our results for weighted average latency social cost. We write OPT to denote the (cost of the) socially optimal outcome, BME for the best mediated equilibrium, and BCE for the best correlated equilibrium. We write n to denote the number of players/jobs and, in weighted games, Δ to denote the ratio of the total weight of jobs to the weight of the smallest job. Note that $(1 + \sqrt{2})/2 \approx 1.2701$.

Related work. Koutsoupias and Papadimitriou initiated the study of the price of anarchy in load-balancing games, considering weighted players, linear latencies, and the maximum (rather than average) social-cost function [24]. A substantial body of follow-up work improved and generalized their initial results [10, 11, 14, 26]. See [18] and [36] for surveys. A second line of work takes social cost to be the sum of players’ costs. Lücking et al. [25] and Gairing et al. [19] measure the price of anarchy of mixed equilibria in linear and convex routing games in this setting. Awerbuch et al. [5] consider both the unweighted and weighted cases on general networks. Kothari et al. [23] and Suri et al. [34] examine the effects of asymmetry in these games (certain links may be unavailable to certain players). Caragiannis et al. [8] give improved bounds on the price of anarchy and stability.

There is also a good deal of work on the quality of correlated equilibria in these games. Christodoulou and Koutsoupias [9, 10] bound the best- and worst-case correlated equilibria in addition to improving existing price of anarchy and stability results. Ashlagi, Monderer, and Tennenholtz [2] introduce and analyze the “mediation value” BCE/BNE and the “enforcement value” OPT/BCE in games with positive payoffs, including a class of games that generalizes load balancing. (Note that the present work considers costs, and we study the reciprocal of their ratios.) Ashlagi et al. also point out a disconnect between the analysis of cost-based and payoff-based games in this setting; one cannot immediately use a payoff-based bound to infer bounds in a cost-based setting, or vice versa. Other aspects of correlated equilibria have also been explored recently, including their existence [20] and computation [20, 21, 28].

Mediated equilibria developed in the game theory literature over time; see Tennenholtz [35] for a summary. Mediated equilibria are studied for position auctions [1], for network routing games [31, 32], and in the context of social choice and voting [29, 30]. Strong mediated equilibria (those resistant to group deviations) are also considered [27, 32].

Some work on mediated equilibria and so-called “delegation equilibria” also explores “Folk Theorem-like” results for one-shot games. Examples include the work of Monderer and Tennenholtz [27], Fershtman et al. [13], and Kalai et al. [22]. A paper by Borgs et al. on the “myth of the Folk Theorem” [6] addresses issues of computational complexity in finding equilibria in repeated games.

2 Notation and Background

Load-balancing games. An n -player, m -machine load-balancing game Γ is defined by a nondecreasing latency function $f_j: [0, \infty) \rightarrow [0, \infty]$ for each machine $j \in \{1, \dots, m\}$, and a weight $w_i > 0$ for each player $i \in \{1, \dots, n\}$. We consider games in which every job has access to every machine: a pure strategy profile $\mathbf{s} = \langle s_1, \dots, s_n \rangle$ can be any element of $\mathcal{S} := \{1, \dots, m\}^n$. The load ℓ_j on a machine j under \mathbf{s} is $\sum_{i: s_i=j} w_i$, and the latency of machine j is $f_j(\ell_j)$. The cost $c_i(\mathbf{s})$ to player i under \mathbf{s} is $f_{s_i}(\ell_{s_i})$. Pure Nash equilibria exist in all load-balancing games [12, 15] (or see [36] or Theorem 5.1). A load-balancing game is *linear* if each f_j is of the form $f_j(x) = a_j \cdot x$ for some $a_j \geq 0$ and *unweighted* if each $w_i = 1$.

Machine j is (strictly) *dominated* by machine j' for player i if, no matter what machines the other $n - 1$ players use, player i 's cost is (strictly) lower using machine j' than using machine j .

Mediators and mediated equilibria. A nonempty subset of the players is called a *coalition*. A *mediator* is a collection Ψ of probability distributions, one for each coalition T , where the probability distribution $\psi_T \in \Psi$ is over pure strategy profiles for the players in T . The *mediated game* M_Γ^Ψ is a new n -player game in which every player either participates in Γ directly by choosing a machine in $S := \{1, \dots, m\}$ or participates by *delegating*. That is, the set of pure strategies in M_Γ^Ψ is $Z = S \cup \{s_{\text{med}}\}$. If the set of delegating players is T , then the mediator plays the correlated strategy ψ_T on behalf of the members of T . In other words, for a strategy profile $\mathbf{z} = \langle z_1, z_2, \dots, z_n \rangle$ where $T := \{i : z_i = s_{\text{med}}\}$ and $\bar{T} := \{i : z_i \neq s_{\text{med}}\} = \{i : z_i \in S\} = \{1, \dots, n\} - T$, the mediator chooses a strategy profile \mathbf{s}_T according to the distribution ψ_T , and plays s_i on behalf of every player $i \in T$; meanwhile, each player i in \bar{T} simply plays z_i . The expected cost to player i under the strategy profile \mathbf{z} is then given by $c_i(\mathbf{z}) := \sum_{\mathbf{s}_T} c_i(\mathbf{s}_T, \mathbf{z}_{\bar{T}}) \cdot \psi_T(\mathbf{s}_T)$. (The mediators described here are called *minimal mediators* in [27], in contrast to a seemingly richer class of mediators—though in fact equally powerful in complete-information games—that allow more communication from players to the mediator.)

A *mediated equilibrium* for Γ is a mediator Ψ such that the strategy profile $\langle s_{\text{med}}, s_{\text{med}}, \dots, s_{\text{med}} \rangle$ is a pure Nash equilibrium in M_Γ^Ψ . Every probability distribution ψ' over the set of all pure strategy profiles for Γ naturally corresponds to a mediator Ψ , where the probability distribution ψ_T for a coalition T is the marginal distribution for T under ψ' —that is, $\psi_T(\mathbf{s}_T) = \sum_{\mathbf{s}' : \mathbf{s}'_T = \mathbf{s}_T} \psi'(\mathbf{s}')$. If ψ' is a correlated equilibrium then this Ψ is a mediated equilibrium.

Social cost and the price of anarchy, stability, etc. The *social cost* of a strategy profile \mathbf{s} is the total (or, equivalently, average) cost of the jobs under \mathbf{s} , weighted by their sizes—that is, $\sum_i w_i \cdot c_i(\mathbf{s})$. (We discuss other social-cost functions in Section 6.)

For a probability distribution ψ over pure strategy profiles, the social cost of ψ is the expected social cost for \mathbf{s} drawn from ψ . We denote by OPT the (cost of the) profile \mathbf{s} that minimizes the social cost, and refer to OPT as *socially optimal*. We denote the worst Nash equilibrium—the one that maximizes social cost—by WNE, and the best Nash equilibrium by BNE. Similarly, we denote the best correlated and best mediated equilibria by BCE and BME, respectively. Note that

$$\text{OPT} \leq \text{BME} \leq \text{BCE} \leq \text{BNE} \leq \text{WNE} \tag{2.1}$$

because every Nash equilibrium is a correlated equilibrium, etc. The *price of anarchy* is WNE/OPT , and the *price of stability* is BNE/OPT . When we refer to the price of stability/anarchy of a class of games, we mean the maximum (or supremum) price of stability/anarchy among all games in the class.

3 Unweighted Load-Balancing Games

Although the unweighted case turns out to have less interesting texture than the weighted version, we start with it because it is simpler and allows us to develop some intuition about the problem. We begin with an illustrative example:

Example 3.1. There are n unweighted jobs and two machines L and R with latency functions $f_L(x) = 1 + \varepsilon$ for any load, and $f_R(x) = 1$ for load $x > n - 1$ and $f_R(x) = 0$ otherwise.

For each player, strategy R strictly dominates strategy L , so the strategy profile $\langle R, R, \dots, R \rangle$ is the unique correlated and Nash equilibrium, with social cost n . Consider the following mediator Ψ . When all n players delegate, the mediator picks uniformly at random from the n strategy profiles in which exactly one player is assigned to L and the remaining $n - 1$ players are assigned to R . When any other set of players delegates, those players are deterministically assigned to R . If all players delegate under Ψ , then each player's expected cost is $(1 + \varepsilon)/n$; if any player deviates, then that player will incur cost at least 1. Thus Ψ is a mediated equilibrium. Its social cost is only $1 + \varepsilon$, which is optimal, while $\text{BNE} = \text{BCE} = n$.

In fact, this “randomize among social optima” technique generalizes to all unweighted load-balancing games: in any such game, $\text{BME} = \text{OPT}$. This is a special case of a general theorem of Monderer and Tennenholtz [27] about mediated equilibria robust to deviations by coalitions. (See also [32].) For completeness, we present here a self-contained proof of this special case of Monderer and Tennenholtz's theorem, showing that $\text{BME} = \text{OPT}$ for any symmetric game. An n -player game is *symmetric* if there is a single shared strategy set S from which each player selects a strategy, and for any permutation π of $\langle 1, \dots, n \rangle$ and any strategy profile \mathbf{s} , we have that the cost c_i incurred by every player i satisfies $c_i(\mathbf{s}) = c_{\pi(i)}(\pi(\mathbf{s}))$. Note that all unweighted load-balancing games are symmetric, as are all single-source single-sink routing games, among many others.

Theorem 3.2 (Monderer and Tennenholtz [27]). *In any symmetric game, we have $\text{BME} = \text{OPT}$.*

Proof. Let \mathbf{o} denote any socially optimal outcome in the game, and let \mathbf{q} denote any Nash equilibrium in the game. Because the game is symmetric, without loss of generality we can permute the players so that player n is the one whose cost is the highest under \mathbf{q} , so that $\text{argmax}_i c_i(\mathbf{q}) = n$. We now define a mediator that behaves as follows:

- If all players delegate, then the mediator chooses a permutation π of $\langle 1, \dots, n \rangle$ uniformly at random and has them play according to $\pi(\mathbf{o})$. In other words, the mediator has all n delegating players play as each player in \mathbf{o} a “fair share” of the time. This permutation is possible because the game is symmetric. If all players delegate, then the expected cost will be equal for all players, and therefore will be a “fair share”— $1/n$ —of the social cost of \mathbf{o} .
- If one player i does not delegate (and the other $n - 1$ players do), then the mediator has players $\langle 1, \dots, i - 1, i + 1, \dots, n \rangle$ play the strategies $\mathbf{q}(1, \dots, n - 1)$, respectively. In other words, the

mediator has all $n - 1$ delegating players play as the first $n - 1$ players in the Nash equilibrium \mathbf{q} , where the n th player in \mathbf{q} is the one who incurs the highest cost.

- Technically, the mediator must be specified even if fewer than $n - 1$ players delegate, so we specify that otherwise the mediator has all delegating players play arbitrarily.

We claim that this mediator defines a mediated equilibrium. For any player i and any strategy s to which player i might deviate from the mediator,

$$\begin{aligned}
 & i\text{'s cost by deviating from the mediator to strategy } s \\
 &= \text{cost to player } n \text{ in playing } s \text{ against } \mathbf{q}(1, \dots, n - 1) \\
 &\geq \text{cost to player } n \text{ in playing } \mathbf{q}(n) \text{ against } \mathbf{q}(1, \dots, n - 1) \\
 &\quad (\mathbf{q} \text{ is a Nash equilibrium, so } \mathbf{q}(n) \text{ is a best response to } \mathbf{q}(1, \dots, n - 1)) \\
 &\geq (1/n) \cdot \text{social cost of } \mathbf{q} \quad (n\text{'s cost under } \mathbf{q} \text{ is maximum among players; max cost } \geq \text{average cost}) \\
 &\geq (1/n) \cdot \text{social cost of } \mathbf{o} \quad (\mathbf{o} \text{ is socially optimal, so } \mathbf{o} \text{ has lower social cost than } \mathbf{q}) \\
 &= i\text{'s cost by delegating,}
 \end{aligned}$$

where the last equality follows because player i receives a “fair share” of the social optimum because i has an equal probability of being each of the players in \mathbf{o} . Thus a player’s cost can only increase by deviating from the mediator, and the mediator defines a mediated equilibrium. When all players delegate, then the only outcome that is ever played is \mathbf{o} (in all permutations), and thus the social cost of this mediator is precisely the social cost of \mathbf{o} , which is OPT by the definition of \mathbf{o} . \square

Example 3.1 shows that with nonlinear latency functions BCE may be much worse than OPT, even in the unweighted 2-machine case. But even linear unweighted load balancing has a gap between BCE and OPT, even in the 2-job, 2-machine case. An example due to Caragiannis et al. [8] demonstrates this gap:

Example 3.3 (Caragiannis et al. [8]). There are two unweighted jobs and there are two machines L and R with latency functions $f_L(x) = x$ and $f_R(x) = (2 + \varepsilon) \cdot x$.

It is easy to verify that BCE and OPT incur costs of 4 and $3 + \varepsilon$, respectively, in this example. (Here machine L is a strictly dominant strategy, so no player can be induced to use machine R in any correlated equilibrium.) Example 3.3 is tight with respect to the gap between BME and BCE, which is a simple consequence of the result of Caragiannis et al. on the price of stability in linear unweighted load-balancing games [8]:

Lemma 3.4. *In linear unweighted load balancing, we have $\text{BCE} \leq \frac{4}{3} \cdot \text{BME}$, and this bound is tight.*

Proof. Caragiannis et al. [8] prove that the price of stability—that is, the ratio between BNE and OPT—is at most $4/3$ in unweighted load-balancing games with linear latency functions, and that Example 3.3 is a tight example for the price of stability. Our lemma follows by (2.1): we have $\text{BCE} \leq \text{BNE} \leq \frac{4}{3} \cdot \text{OPT} \leq \frac{4}{3} \cdot \text{BME}$, and Example 3.3 remains tight for the ratio BCE/BME , because $\text{OPT} = \text{BME}$ by Theorem 3.2. \square

Assembling Lemma 3.4 and Theorem 3.2 yields a complete picture for unweighted linear-latency load-balancing games. We can also show a tight bound for the ratio BCE/BME in unweighted *nonlinear* latency load-balancing games:

Lemma 3.5. *In any n -player unweighted load-balancing game with not-necessarily-linear latency functions, we have $\text{BCE} \leq n \cdot \text{BME}$, and this bound is tight.*

Proof. To establish that $\text{BCE} \leq n \cdot \text{BME}$, we prove a slightly stronger claim, namely

(\dagger) the cost of the worst pure Nash equilibrium WPNE satisfies $\text{WPNE} \leq n \cdot \text{OPT}$.

We first show how the lemma follows from (\dagger). Because every pure Nash equilibrium is a correlated equilibrium (so the best correlated equilibrium's cost cannot exceed that of any pure Nash equilibrium, including WPNE), we have $\text{BCE} \leq \text{WPNE}$. Also, we have $\text{OPT} = \text{BME}$ by Theorem 3.2. Thus $\text{BCE} \leq \text{WPNE} \leq n \cdot \text{OPT} = n \cdot \text{BME}$. Note that for this argument to go through it is crucial that at least one pure Nash equilibrium must exist, which, as discussed in Section 2, is true in all load-balancing games.

We now turn to the proof of (\dagger). Let \mathbf{q} be an arbitrary pure Nash equilibrium, and fix any player i . We claim that player i 's cost under \mathbf{q} is at most OPT . For a machine j , let $k_j(\mathbf{q})$ and $k_j(\text{OPT})$ denote the number of jobs assigned to machine j by \mathbf{q} and OPT , respectively. If $\mathbf{q} = \text{OPT}$, we are done; otherwise $\mathbf{q} \neq \text{OPT}$, and there exists a machine j^* such that $k_{j^*}(\mathbf{q}) < k_{j^*}(\text{OPT})$, because $\sum_j k_j(\mathbf{q}) = \sum_j k_j(\text{OPT})$. Furthermore, because each k_j is integral for pure strategy profiles like \mathbf{q} and OPT , we have $k_{j^*}(\mathbf{q}) + 1 \leq k_{j^*}(\text{OPT})$. Thus

$$\begin{aligned}
 & i\text{'s cost under } \mathbf{q} \\
 &= i\text{'s cost for using machine } \mathbf{q}(i) \text{ if others play according to } \mathbf{q} \\
 &\leq i\text{'s cost for using machine } j^* \text{ if others play according to } \mathbf{q} \\
 &\quad (\mathbf{q} \text{ is a Nash equilibrium, so } \mathbf{q}(i) \text{ is a best response for player } i) \\
 &\leq \text{cost of one job using } j^* \text{ under load } k_{j^*}(\mathbf{q}) + 1 \\
 &\quad (\text{the load of } j^* \text{ is one for player } i \text{ plus at most } k_{j^*}(\mathbf{q}) \text{ for the other players}) \\
 &\leq \text{cost of one job using } j^* \text{ under load } k_{j^*}(\text{OPT}) \\
 &\quad (k_{j^*}(\mathbf{q}) + 1 \leq k_{j^*}(\text{OPT}) \text{ as above; machine costs are nondecreasing with load}) \\
 &\leq \text{the total cost of } \text{OPT}.
 \end{aligned}$$

Thus each player's cost in \mathbf{q} is at most OPT , and the total cost of \mathbf{q} is at most $n \cdot \text{OPT}$.

Although the last inequality seems to be giving away a huge amount by lower bounding the total cost of OPT by the cost incurred by a single player under OPT , in fact this bound is tight: Example 3.1 was an n -player unweighted load-balancing game in which $\text{BCE} = \text{BNE} = n$, but $\text{OPT} = \text{BME} = 1 + \varepsilon$. \square

This lemma now completes the picture for all unweighted load-balancing games:

Theorem 3.6. *In n -player unweighted load-balancing games:*

- (i) *with linear latency functions, $\text{BCE} \leq \frac{4}{3} \cdot \text{BME}$ and $\text{BME} = \text{OPT}$.*
- (ii) *with not-necessarily-linear latency functions, $\text{BCE} \leq n \cdot \text{BME}$ and $\text{BME} = \text{OPT}$.*

All of these bounds are tight.

Proof. The result follows immediately from Theorem 3.2, Theorem 3.4, and Theorem 3.5. \square

4 Weighted Linear Load-Balancing Games

We now turn to weighted load-balancing games, where we find a richer landscape of results: among other things, we will illustrate cases in which the best mediated equilibrium falls strictly between the best correlated equilibrium and the social optimum. We begin with the linear-latency case, focusing on 2-player games. We then provide bounds for general games with linear latencies. We begin by proving the following theorem:

Theorem 4.1. *In any 2-machine, 2-player weighted load-balancing game with linear latency functions:*

1. $\text{BCE}/\text{BME} \leq 4/3$. This bound is tight for an instance with weights $\{1, 1\}$ and with latency functions $f_L(x) = x$ and $f_R(x) = (2 + \epsilon) \cdot x$.
2. $\text{BME}/\text{OPT} \leq \frac{1+\sqrt{2}}{2}$. This bound is tight for an instance with weights $\{1, 1 + \sqrt{2}\}$ and with latency functions $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x$.

Proof. The theorem follows immediately from Theorem 4.9 and Theorem 4.11 below. □

This result fully resolves the 2-player, 2-machine case with linear latency functions. The worst case for BCE/BME is actually unweighted—in fact, Example 3.3—but there is a weighted instance in which the best mediated equilibrium incurs a social cost of $\frac{1+\sqrt{2}}{2} \cdot \text{OPT} \approx 1.2071 \cdot \text{OPT}$; furthermore, this bound is tight for 2-player games. Also included in this class of games are instances in which $\text{BCE} > \text{BME} > \text{OPT}$. One concrete example is with weights $\{1, 1 + \sqrt{2}\}$, $f_L(x) = x$, and $f_R(x) = \frac{3+3\sqrt{2}}{2} \cdot x$, when $\text{BCE}/\text{BME} = \frac{20+4\sqrt{2}}{23} \approx 1.1155$ and $\text{BME}/\text{OPT} = \frac{14\sqrt{2}-1}{17} \approx 1.1058$.

Adding additional machines to a 2-player instance does not substantively change the results (there is no point in either player using anything other than the two “best” machines), but the setting with $n \geq 3$ players requires further analysis, and, it appears, new techniques. Recent results on the price of anarchy in linear load-balancing games [5, 8, 10] imply an upper bound of $1 + \phi \approx 2.618$ on BME/OPT for any number of players n , where ϕ is the golden ratio. We can improve this upper bound to 2 by bounding the price of stability using results communicated to us by Ioannis Caragiannis [7]; see Theorem 4.12. We believe that the worst-case ratio of BME/OPT does not decrease as n increases (for example, we can consider n -player instances in which $n - 2$ players have jobs of negligible weight and the remaining two players have jobs as in Theorem 4.1), but we do not have a proof that things cannot get worse as n increases, nor do we have a 3-job example for which the best mediated equilibrium is worse than $\frac{1+\sqrt{2}}{2} \cdot \text{OPT} \approx 1.2071 \cdot \text{OPT}$. The major open challenge emanating from our work is to close the gap between the upper bound ($\text{BME}/\text{OPT} \leq 2$) and our bad example ($\text{BME}/\text{OPT} = 1.2071$) for general n .

4.1 Analysis of 2-Machine, 2-Player Weighted Linear-Latency Load Balancing

Before embarking on the proof, we fix some notation for the remainder of the section. We have two machines (“left” and “right”, abbreviated L and R), with latency functions $f_L(x) = x$ and $f_R(x) = ax$, for $a \geq 1$, respectively. We have two jobs (“little” and “big”), with weights 1 and $w \geq 1$, respectively.

Writing this load-balancing game in normal form yields the following bimatrix of costs:

		<i>“big” player</i>	
		<i>L</i>	<i>R</i>
<i>“little” player</i>	<i>L</i>	$w + 1, w + 1$	$1, aw$
	<i>R</i>	a, w	$a(w + 1), a(w + 1)$.

We abbreviate strategy profiles by writing the little player’s strategy followed by the big player’s strategy; for example, LR denotes the little player’s use of machine L and the big player’s use of machine R . Recall that the social-cost function is the weighted total cost: the cost to the little player plus w times the cost to the big player.

Lemma 4.2. *The social optimum has cost*

$$\text{OPT} = \begin{cases} (w + 1)^2 & \text{if } a \geq 2w + 1 \\ a + w^2 & \text{otherwise.} \end{cases}$$

Proof. We claim that the social optimum is never uniquely achieved by RR or LR . The social cost of RR is $a(w + 1)^2$, which is $a \geq 1$ times the social cost of LL . Similarly, the social cost of RL is $a + w^2 \leq 1 + aw^2$, the social cost of LR . Thus the social optimum is achieved by either LL or RL , at total social cost $(w + 1)^2$ or $a + w^2$. Observe that $(w + 1)^2 = w^2 + 2w + 1$ exceeds $a + w^2$ if and only if $a < 2w + 1$. \square

Lemma 4.3. *The best Nash equilibrium and the best correlated equilibrium have cost*

$$\text{BNE} = \text{BCE} = \begin{cases} (w + 1)^2 & \text{if } a > w + 1 \\ a + w^2 & \text{otherwise.} \end{cases}$$

Proof. Suppose $a > w + 1$. Then L is a strictly dominant strategy for both players, and thus the only Nash equilibrium is for both to play L . Furthermore, no correlated equilibrium can play any outcome other than LL with nonzero probability, because R is a strictly dominated strategy.

On the other hand, suppose $a \leq w + 1$. Then RL is a Nash equilibrium: the little player prefers the cost of a to LL ’s cost of $w + 1$, and the big player certainly prefers w to $a(w + 1)$. The social cost of RL is $a + w^2$. By the inequality $a \leq w + 1 < 2w + 1$ and Theorem 4.2, $a + w^2$ is the social optimum; there cannot be a better Nash equilibrium or correlated equilibrium. \square

Lemma 4.4. *Suppose $a > w + 1$. In the best mediated equilibrium, if one player deviates from the mediator then the mediator will have the one remaining delegating player choose L . Under this mediator, the cost incurred by a deviating player is $1 + w$.*

Proof. The mediator has the power to play a randomized strategy on behalf of the one remaining delegating player. Suppose the mediator has the delegating player play L with probability $p \in [0, 1]$.

If the little player deviates from delegating, he will deviate to whichever machine induces a lower cost for him. Thus his cost of deviating is

$$\min(p(w + 1) + (1 - p), pa + (1 - p)a(w + 1)) = pw + 1,$$

because the cost of deviating to R is at least $a > w + 1 \geq pw + 1$. If, on the other hand, the big player deviates from delegating, then her cost is

$$\min(p(w + 1) + (1 - p)w, paw + (1 - p)a(w + 1)) = w + p.$$

In both cases, the cost to the deviating player is a strictly increasing function of p . Thus the cost to the deviating player is maximized when $p = 1$, when $pw + 1 = w + p = w + 1$. Maximizing the players' costs for deviating from the mediator can never destabilize a mediated equilibrium, so the best mediated equilibrium's mediator might as well use this punishment strategy. \square

It will be useful to fix some notation for the rest of this section. Denote by $\langle\langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle\rangle$ the mediator that (1) has any delegating player play L if the other player deviates, and (2) plays LL , LR , RL , and RR with respective probabilities p_{LL} , p_{LR} , p_{RL} , and p_{RR} if both players delegate. (Note that the mediator $\langle\langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle\rangle$ is defined only when $p_{LL} + p_{LR} + p_{RL} + p_{RR} = 1$.)

Lemma 4.5. *Suppose $a > w + 1$. The mediator $M = \langle\langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle\rangle$ forms a mediated equilibrium if and only if both of the following conditions hold:*

$$w + 1 \geq (w + 1)p_{LL} + p_{LR} + ap_{RL} + a(w + 1)p_{RR}, \quad (4.1)$$

$$w + 1 \geq (w + 1)p_{LL} + awp_{LR} + wp_{RL} + a(w + 1)p_{RR}. \quad (4.2)$$

Furthermore, the social cost of M is the weighted sum of the right-hand sides of (4.1) and (4.2).

Proof. By the definition of the mediator and by Theorem 4.4, the cost to a nondelegating player is $w + 1$. The costs to the two players if both delegate are the right-hand sides of (4.1) and (4.2), respectively. Thus M is a mediated equilibrium if and only if both constraints are satisfied. \square

Lemma 4.6. *Suppose $a > w + 1$. If $M = \langle\langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle\rangle$ is a mediated equilibrium, then so is the mediator $M' = \langle\langle p_{LL} + p_{RR}, p_{LR}, p_{RL}, 0 \rangle\rangle$ that shifts all the probability of playing RR to the outcome LL . Furthermore, the social cost of M' is no larger than the social cost of M .*

Proof. Suppose M is a mediated equilibrium—i.e., by Theorem 4.5, suppose that p_{LL} , p_{LR} , p_{RL} , and p_{RR} satisfy (4.1) and (4.2). Because $a(w + 1) \geq w + 1$, the right-hand sides of both (4.1) and (4.2) have not increased when we move from M to M' , so the social cost has not increased while the constraints remain satisfied. Thus the social cost of M' is no greater than the social cost of M , and M' remains a mediated equilibrium. \square

Lemma 4.7. *Suppose $a > w + 1$. For any two probabilities $p > 0$ and $q \geq 0$, define the mediator $M_{p,q} = \langle\langle 1 - p, pq, p(1 - q), 0 \rangle\rangle$.*

- (i) *If $a > w + 1 + \frac{1}{w}$, then no $M_{p,q}$ is a mediated equilibrium.*
- (ii) *If $a \leq w + 1 + \frac{1}{w}$, then there are values of p, q for which $M_{p,q}$ is a mediated equilibrium. Furthermore, the mediated equilibrium $M_{p,q}$ with the lowest social cost is achieved when $q = (a - w - 1)/(a - 1)$ and $p = 1$, when the cost is $aw^2 - w^3 + w + 1$.*

Proof. Applying Theorem 4.5 to the mediator $M_{p,q} = \langle \langle 1-p, pq, p(1-q), 0 \rangle \rangle$ implies that $M_{p,q}$ is a mediated equilibrium if and only if

$$\begin{aligned} w+1 &\geq (w+1)(1-p) + ap(1-q) + pq, \\ w+1 &\geq (w+1)(1-p) + wp(1-q) + awpq. \end{aligned}$$

Collecting like terms and dividing both inequalities by p , we have

$$w+1 \geq a(1-q) + q, \quad w+1 \geq w(1-q) + awq.$$

Solving for q yields

$$\frac{a-w-1}{a-1} \leq q \leq \frac{1}{(a-1)w}. \quad (4.3)$$

Notice that there is a q satisfying (4.3) if and only if

$$\frac{a-w-1}{a-1} \leq \frac{1}{(a-1)w} \iff a \leq w+1 + \frac{1}{w}.$$

Thus, if $a > w+1 + \frac{1}{w}$, then no $M_{p,q}$ is a mediated equilibrium. For the case in which $a \leq w+1 + \frac{1}{w}$, notice that the social cost of $M_{p,q}$ is given by

$$\begin{aligned} \text{social cost of } M_{p,q} &= (1-p)(w+1)^2 + pq(1+aw^2) + p(1-q)(a+w^2) \\ &= (w+1)^2 + p(a-2w-1) + (a-1)(w^2-1)pq. \end{aligned} \quad (4.4)$$

By assumption, we have $a \geq 1$, $w \geq 1$, and $p > 0$, so this social cost is an increasing function of q . Thus the social cost is minimized for the smallest value of q such that (4.3) is satisfied, namely $q^* = (a-w-1)/(a-1)$. Plugging this value for q^* into (4.4), we have that

$$\begin{aligned} \text{social cost of } M_{p,q^*} &= (w+1)^2 + p(a-2w-1) + (w^2-1)p(a-w-1) \\ &= (w+1)^2 + pw((a-1)w-1-w^2). \end{aligned} \quad (4.5)$$

By assumption, we have that $a \leq w+1 + \frac{1}{w}$, and thus

$$(a-1)w-1-w^2 \leq (w+1 + \frac{1}{w} - 1)w-1-w^2 = (w + \frac{1}{w})w-1-w^2 = 0.$$

Because $w > 0$, the social cost in (4.5) is a decreasing function of p , so the social cost is minimized for the largest value of p , namely $p = 1$. In this case, the social cost is

$$(w+1)^2 + w((a-1)w-1-w^2) = aw^2 - w^3 + w + 1,$$

and we are done. □

Lemma 4.8. *The best mediated equilibrium has cost*

$$\text{BME} = \begin{cases} (w+1)^2 & \text{if } a > w+1 + \frac{1}{w}, \\ aw^2 - w^3 + w + 1 & \text{if } a \in (w+1, w+1 + \frac{1}{w}], \\ a + w^2 & \text{if } a \leq w+1. \end{cases}$$

	$a \in [1, w+1]$	$a \in (w+1, w+1 + \frac{1}{w}]$	$a \in (w+1 + \frac{1}{w}, 2w+1]$	$a \in (2w+1, \infty)$
OPT	$a + w^2$	$a + w^2$	$a + w^2$	$(w+1)^2$
BME	$a + w^2$	$aw^2 - w^3 + w + 1$	$(w+1)^2$	$(w+1)^2$
BCE = BNE	$a + w^2$	$(w+1)^2$	$(w+1)^2$	$(w+1)^2$

Figure 2: Social costs of OPT, BME, BCE, and BNE for the 2-player, 2-machine game with jobs of weight 1 and $w \geq 1$ and with machines with latency functions $f_L(x) = x$ and $f_R(x) = ax$, for $a \geq 1$.

Proof. For $a \leq w+1$, the best Nash equilibrium achieves the social optimum, by Theorem 4.2 and Theorem 4.3. Any Nash equilibrium is a mediated equilibrium, and no mediated equilibrium can outperform the social optimum; thus $\text{BME} = \text{OPT} = a + w^2$.

Henceforth assume that $a > w+1$. By Theorem 4.4, the best mediated equilibrium is achieved by a mediator of the form $\langle\langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle\rangle$. We may assume that $p_{RR} = 0$ by Theorem 4.6. The best mediated equilibrium is then either $\langle\langle 1, 0, 0, 0 \rangle\rangle$ or some $M_{p,q}$. (If $p_{LL} < 1$, then $\langle\langle p_{LL}, p_{LR}, p_{RL}, 0 \rangle\rangle = M_{p,q}$ with $p = 1 - p_{LL}$ and $q = \frac{p_{LR}}{1 - p_{LL}}$.)

For $a > w+1 + \frac{1}{w}$, by Theorem 4.7 there is no mediated equilibrium of the form $M_{p,q}$, so the only mediated equilibrium is the $\langle\langle 1, 0, 0, 0 \rangle\rangle$ mediator, which is in fact the pure Nash equilibrium from Theorem 4.3, with social cost $(w+1)^2$.

For $w+1 < a \leq w+1 + \frac{1}{w}$, the best mediated equilibrium of the form $M_{p,q}$ has social cost $aw^2 - w^3 + w + 1$ by Theorem 4.7. For values of a in this range, we have that $aw^2 - w^3 + w + 1 \leq (w+1)^2$, by the same argument as in the proof of Theorem 4.7, so this mediator achieves a better social cost than the $\langle\langle 1, 0, 0, 0 \rangle\rangle$ mediator. \square

These lemmas characterizing BNE, BCE, BME, and OPT are summarized in Figure 2.

Theorem 4.9. *In any 2-machine, 2-job weighted load-balancing game with linear latency functions, we have $\text{BCE}/\text{BME} \leq 4/3$. Furthermore, this bound is tight for jobs of weights 1 and machines with latency functions $f_L(x) = x$ and $f_R(x) = (2 + \varepsilon) \cdot x$.*

Proof. By Theorem 4.3 and Theorem 4.8, the best correlated and best mediated equilibria have the same cost except when $w+1 < a \leq w+1 + 1/w$, where

$$\text{BCE}/\text{BME} = \frac{(w+1)^2}{aw^2 - w^3 + w + 1}.$$

This ratio is a decreasing function of a , so the ratio approaches its maximum at the smallest value of a in this range, namely as a approaches $w+1$ from above. Thus

$$\text{BCE}/\text{BME} \leq \frac{(w+1)^2}{(w+1)w^2 - w^3 + w + 1} = \frac{w^2 + 2w + 1}{w^2 + w + 1}.$$

This ratio achieves its maximum of $4/3$ when $w = 1$. Thus the worst case is actually an unweighted example—in fact, the same example from Caragiannis et al. [8] and Theorem 3.6. \square

Lemma 4.10. *The ratio between the cost of the best mediated equilibrium and the social optimum is at most $\frac{w^2+w}{w^2+1}$, and that bound is tight when $a = w+1 + 1/w$.*

Proof. By Theorem 4.2 and Theorem 4.8, we have the cost of the social optimum and of the best mediated equilibrium. We will consider the various ranges for a defined by those lemmas.

If $a \leq w + 1$, then both the social optimum and the best mediated equilibrium have social cost $a + w^2$, for a ratio of one.

If $w + 1 < a \leq w + 1 + \frac{1}{w}$, then the best mediated equilibrium has cost $aw^2 - w^3 + w + 1$ and the social optimum has cost $a + w^2$. Thus the ratio is

$$\frac{aw^2 - w^3 + w + 1}{a + w^2} = \frac{aw^2 + w^4}{a + w^2} + \frac{-w^4 - w^3 + w + 1}{a + w^2} = w^2 + \frac{(w + 1)(1 - w^3)}{a + w^2}.$$

Because $w \geq 1$, we have that $(w + 1)(1 - w^3) \leq 0$, and thus this ratio is an increasing function of a . Therefore the ratio throughout this range is upper-bounded by its value at $a = w + 1 + \frac{1}{w}$, when its value is

$$\frac{(w + 1 + \frac{1}{w})w^2 - w^3 + w + 1}{w + 1 + \frac{1}{w} + w^2} = \frac{w^2 + 2w + 1}{(w + 1)(w + \frac{1}{w})} = \frac{w + 1}{w + \frac{1}{w}} = \frac{w^2 + w}{w^2 + 1}. \quad (4.6)$$

If $w + 1 + \frac{1}{w} < a < 2w + 1$, then the best mediated equilibrium has cost $(w + 1)^2$ and the social optimum has cost $a + w^2$. Thus the ratio of their costs is $\frac{(w+1)^2}{a+w^2}$, which is a decreasing function of a . Therefore the ratio throughout this range is upper-bounded by its value at $a = w + 1 + \frac{1}{w}$. For this value of a , we have $aw^2 - w^3 + w + 1 = (w + 1)^2$, and thus the ratio is upper-bounded throughout this range by $\frac{w^2+w}{w^2+1}$, just as in (4.6).

If $a \geq 2w + 1$, then both the social optimum and the best mediated equilibrium have social cost $(w + 1)^2$, for a ratio of one. \square

Theorem 4.11. *In any 2-machine, 2-job weighted load-balancing game with linear latency functions, we have $\text{BME}/\text{OPT} \leq \frac{1+\sqrt{2}}{2} \approx 1.2071$. Furthermore, this bound is tight for jobs of weights 1 and $1 + \sqrt{2}$ and machines with latency functions $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x$.*

Proof. Without loss of generality, we scale the jobs and latency functions so that the jobs' weights are 1 and $w \geq 1$ and the machines' latency functions are $f_L(x) = x$ and $f_R(x) = ax$ for $a \geq 1$. By Theorem 4.10, the worst-case ratio of the best mediated equilibrium to the social optimal outcome occurs when $a = w + 1 + 1/w$, when

$$\frac{\text{BME}}{\text{OPT}} = \frac{w^2 + w}{w^2 + 1}.$$

We wish to upper-bound this ratio over all $w \geq 1$. By simple calculus, the ratio is maximized when $w^2 - 2w - 1 = 0$, or when $w = 1 + \sqrt{2}$. For this value of w , we have $\text{BME}/\text{OPT} = \frac{4+3\sqrt{2}}{4+2\sqrt{2}} = \frac{1+\sqrt{2}}{2}$. Thus the ratio of the best mediated equilibrium to the social optimum never exceeds $\frac{1+\sqrt{2}}{2}$.

Furthermore, for the load-balancing game with jobs of weights 1 and $1 + \sqrt{2}$, and with machines with latency functions $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x \approx 3.8284 \cdot x$, the ratio of the best mediated equilibrium to the social optimum is $\frac{1+\sqrt{2}}{2} \approx 1.2071$. \square

4.2 Bound for n -Player, m -Machine Weighted Linear-Latency Load Balancing

We now turn to the general n -player, m -machine case for weighted linear-latency load balancing. In this section, we present a theorem communicated to us by Ioannis Caragiannis [7] to provide an upper bound on the quality of the best Nash equilibrium (and thus bound the quality of the best mediator). The proof relies on a potential function due to Fotakis, Kontogiannis, and Spirakis [14].

Theorem 4.12 (Caragiannis [7]). *For any weighted linear-latency load-balancing game with n players, we have that $\text{BNE} \leq 2 \cdot \text{OPT}$, $\text{BME} \leq 2 \cdot \text{OPT}$, and $\text{BCE} \leq 2 \cdot \text{BME}$.*

Proof. The theorem holds for slightly more general latency functions—of the form $f(x) = ax + b$ instead of just of the form $f(x) = ax$ —so we present the more general result.

Consider any n -job, m -machine linear-latency load-balancing game, and let $f_j(x) = a_jx + b_j$ denote the latency function of machine $j \in \{1, \dots, m\}$. For a pure strategy profile \mathbf{s} , let $j_s = \{i : s_i = j\}$ denote the set of jobs that use machine j under \mathbf{s} . Define the potential of a strategy profile \mathbf{s} as

$$\Phi(\mathbf{s}) := \sum_{j=1}^m \left[a_j \sum_{\substack{i, i' \in j_s \\ i \leq i'}} w_i w_{i'} + b_j \sum_{i \in j_s} w_i \right].$$

We first claim that Φ is in fact a potential function. Let \mathbf{s} be any strategy profile, and let \mathbf{s}' be the strategy profile that results from player i moving his job from machine $j = s_i$ to machine $j' = s'_i$. Then in moving from \mathbf{s} to \mathbf{s}' , when job i is removed from machine j then the value of Φ decreases by

$$a_j \sum_{i' \in j_s} w_i w_{i'} + b_j w_i = w_i \left(a_j \sum_{i' \in j_s} w_{i'} + b_j \right).$$

This quantity is exactly w_i times the cost that i had been paying to use j . Similarly, when i places his job on the other machine j' , the increase in Φ is exactly i 's new cost scaled by w_i . Therefore $\Phi(\mathbf{s}') - \Phi(\mathbf{s})$ is precisely the difference in cost experienced by player i between \mathbf{s}' and \mathbf{s} , scaled by w_i . Thus Φ tracks player improvements scaled by their weights, and as such Φ is a potential function.

Now observe that $\Phi(\mathbf{s})$ is very closely related to the social cost of \mathbf{s} , which we will denote as $c(\mathbf{s})$. In particular,

$$c(\mathbf{s}) = \sum_{j=1}^m \left[a_j \left(\sum_{i \in j_s} w_i \right)^2 + b_j \sum_{i \in j_s} w_i \right] = \sum_{j=1}^m \left[a_j \sum_{i, i' \in j_s} w_i w_{i'} + b_j \sum_{i \in j_s} w_i \right].$$

The only difference between $c(\mathbf{s})$ and $\Phi(\mathbf{s})$ is that in $c(\mathbf{s})$, the coefficient of each $a_j w_i w_{i'}$ term is 2 for any $i \neq i'$, whereas in $\Phi(\mathbf{s})$ the coefficient is 1. Thus we have that

$$\Phi(\mathbf{s}) \leq c(\mathbf{s}) \leq 2 \cdot \Phi(\mathbf{s}) \text{ for any strategy profile } \mathbf{s}. \quad (4.7)$$

Consider any optimal solution \mathbf{o} and let \mathbf{q} be a Nash equilibrium reached by running best-response dynamics (BRD) starting from \mathbf{o} : while there exists a player who can improve her cost by swapping to a different machine, update \mathbf{o} by applying this swap. Plugging \mathbf{q} and \mathbf{o} in for \mathbf{s} in (4.7) gives us

that $\Phi(\mathbf{o}) \leq c(\mathbf{o})$ and $c(\mathbf{q}) \leq 2 \cdot \Phi(\mathbf{q})$. Furthermore, because \mathbf{q} was reached via best-response dynamics from \mathbf{o} , and Φ is a potential function, we have that $\Phi(\mathbf{q}) \leq \Phi(\mathbf{o})$. Stringing these three inequalities together gives us

$$c(\mathbf{q}) \leq 2 \cdot \Phi(\mathbf{q}) \leq 2 \cdot \Phi(\mathbf{o}) \leq 2 \cdot c(\mathbf{o}).$$

The best Nash equilibrium can only be cheaper than \mathbf{q} , and thus we have $\text{BNE} \leq 2 \cdot \text{OPT}$. Because $\text{OPT} \leq \text{BME} \leq \text{BCE} \leq \text{BNE}$, this inequality implies both $\text{BME} \leq 2 \cdot \text{OPT}$ and $\text{BCE} \leq 2 \cdot \text{BME}$, as desired. \square

5 Weighted Nonlinear Load-Balancing Games

We now consider weighted load-balancing games with latency functions that are not necessarily linear. We know from Lemma 3.6 that even in unweighted cases the power of Nash and correlated equilibria is limited. The weighted setting is even worse: the price of anarchy is unbounded, even if we restrict our attention to *pure* equilibria. Consider two identical machines, with latencies $f(x) = 0$ for $x \leq 5$ and $f(x) = 1$ for $x \geq 6$. There are four jobs, two of size 3 and two of size 2. A solution with cost zero exists (each machine has one size-2 and one size-3 job), but putting the two size-3 jobs on one machine and the two size-2 jobs on the other is a pure Nash equilibrium too. We can show that the price of stability is better in this setting, but in general BME is no better than BNE:

Theorem 5.1. *In any n -player weighted load-balancing game with job weights $\{w_1, \dots, w_n\}$ (and not necessarily linear latency functions), $\text{BNE} \leq \Delta \cdot \text{OPT}$, where $\Delta := \sum_i w_i / \min_i w_i$ is the ratio of total job weight to smallest job weight. Thus $\text{BME} \leq \Delta \cdot \text{OPT}$ and $\text{BCE} \leq \Delta \cdot \text{BME}$. Both bounds are tight.*

Proof. Our proof is a special case of the argument that pure Nash equilibria always exist in load-balancing games [12, 15, 36]. Start with $\mathbf{s} := \text{OPT}$, the socially optimal outcome, and run best-response dynamics (BRD): while there exists a player who can improve her cost by swapping to a different machine, update \mathbf{s} by applying this swap. Consider the vector \mathbf{v} of costs experienced by each player. The proof that a pure Nash equilibrium exists argues that the sorted vector \mathbf{v} decreases lexicographically at every iteration, and thus that BRD eventually terminates at a Nash equilibrium [12, 15, 36]. Thus the *maximum* entry in \mathbf{v} can only decrease during each iteration of BRD. Therefore every entry in the final vector \mathbf{v}_f is upper bounded by the maximum entry x in \mathbf{v}_0 , the original vector \mathbf{v} from OPT. The cost of OPT must be at least $x \cdot \min_i w_i$ (if \mathbf{v}_0 has only one non-zero entry which corresponds to the lightest job); the cost of the pure Nash equilibrium found by BRD can be at most $x \cdot \sum_i w_i$ (if every entry of \mathbf{v}_f is equal to x). Thus $\text{BNE} \leq x \cdot \sum_i w_i$ and $\text{OPT} \geq x \cdot \min_i w_i$, and $\text{BNE}/\text{OPT} \leq \Delta$.

Combining $\text{BNE}/\text{OPT} \leq \Delta$ with the facts that $\text{BME} \leq \text{BNE}$, $\text{BCE} \leq \text{BNE}$, and $\text{OPT} \leq \text{BME}$ yields the stated bounds. We have actually already given a tight example for the $\text{BCE} \leq \Delta \cdot \text{BME}$ bound: the unweighted n -job Example 3.1, in which $\text{BCE} = (1 - \varepsilon) \cdot n \cdot \text{BME}$, has $\Delta = n$. To show that the $\text{BME} \leq \Delta \cdot \text{OPT}$ bound is tight, consider the following 2-machine, n -job example. Machine L has zero cost if its load does not exceed $n - 1$ and has unit cost for loads $> n - 1$. Machine R has cost $1 + \varepsilon$ if its load does not exceed ε and has cost ∞ for loads $> \varepsilon$. (A sufficiently large finite cost would suffice, but for simplicity of presentation we use infinite latency.) There are $n - 1$ unit-size jobs and one “small job” of size ε . Any mediated equilibrium must deterministically place all unit-sized jobs on L , because no unit-sized player is willing to incur the infinite cost of R with nonzero probability.

With these $n - 1$ jobs on machine L , the small job is guaranteed to incur a cost of at most 1, even without mediation. Therefore, a mediated equilibrium cannot place the small job on R with a nonzero probability, as the expected cost to the small job would be > 1 . Therefore the only mediated equilibrium puts all n jobs on L deterministically—and thus $\text{BME} = n - 1 + \varepsilon$. The socially optimal solution, on the other hand, puts the small job on R and the other $n - 1$ jobs on L , for a cost of $\text{OPT} = \varepsilon \cdot (1 + \varepsilon)$. Thus $\text{BME}/\text{OPT} = \frac{n-1+\varepsilon}{\varepsilon \cdot (1+\varepsilon)} = \frac{\Delta}{1+\varepsilon}$. \square

6 Other Social-Cost Functions

Two other social-cost functions receive attention in the literature on load-balancing games: the maximum latency of any job, and the unweighted average latency of the jobs.

$$\text{sc}_{\text{wavg}}(\mathbf{s}) := \sum_i w_i \cdot c_i(\mathbf{s}) \quad \text{sc}_{\text{uavg}}(\mathbf{s}) := \sum_i c_i(\mathbf{s}) \quad \text{sc}_{\text{max}}(\mathbf{s}) := \max_i c_i(\mathbf{s})$$

Until now, we have discussed sc_{wavg} exclusively; here, we briefly discuss the other two functions.

Under sc_{max} , in any load-balancing game $\text{BNE}/\text{OPT} = 1$, and thus $\text{BME} = \text{OPT} = \text{BNE}$. The argument is just as in Theorem 5.1: start from OPT and run BRD until it converges; the maximum load on a machine does not increase throughout this process, so the resulting Nash equilibrium remains socially optimal under sc_{max} . Thus mediation is uninteresting under this social-cost function.

Load-balancing games measured by unweighted total latency turn out to have behavior that is mostly qualitatively similar to weighted total latency. Obviously sc_{wavg} and sc_{uavg} are identical if the jobs are unweighted, so only the weighted-job case is interesting. If the latency functions are nonlinear, then an analogue to Theorem 5.1 establishes that $\text{BCE} \leq n \cdot \text{BME}$ and $\text{BME} \leq n \cdot \text{OPT}$ (the total cost of OPT is at least the maximum cost x experienced by a job in OPT ; running BRD from OPT yields a Nash equilibrium where each job experiences cost at most x); the examples from Theorem 5.1 and Example 3.1 both remain tight. The 2-job linear case is also qualitatively similar:

Theorem 6.1. *Under the social-cost function sc_{uavg} , in linear weighted 2-machine 2-job load-balancing games:*

- $\text{BME}/\text{OPT} \leq \frac{2+4\sqrt{2}}{7} \approx 1.0938$. This bound is tight for the example from Theorem 4.1.2: the jobs have weights $\{1, 1 + \sqrt{2}\}$ and the latency functions are $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x$.
- $\text{BCE}/\text{BME} \leq \frac{4}{3}$. The bound is tight for Example 3.3.

Proof. To prove this theorem, we parallel the analysis of Section 4; algebraic details are different, but the analysis goes through using the same overall approach. We derive the following results:

	$a \in [1, w + 1]$	$a \in (w + 1, w + 1 + \frac{1}{w}]$	$a \in (w + 1 + \frac{1}{w}, 2 + w]$	$a \in (2 + w, \infty)$
OPT	$a + w$	$a + w$	$a + w$	$2w + 2$
BME	$a + w$	$aw - w^2 + w + 1$	$2w + 2$	$2w + 2$
BCE = BNE	$a + w$	$2w + 2$	$2w + 2$	$2w + 2$

(The proofs for this social-cost function are found in Appendix A.) The worst-case value of a for BME/OPT is again $a = 1 + w + 1/w$, when $\text{BME}/\text{OPT} = \frac{2w^2+2w}{2w^2+w+1}$. The ratio is maximized when $w = 1 + \sqrt{2}$. The

worst-case value of a for BCE/BME is again $a = 1 + w + \varepsilon$, when $\text{BCE/BME} = \frac{2w+2}{2w+1+\varepsilon}$, which is again maximized at $w = 1$ when the ratio is $4/3$. \square

Intriguingly, the worst-case values for both ratios (BME/OPT and BCE/BME) under sc_{uavg} are achieved at exactly the same values of a and w as under sc_{wavg} (though the numerical value taken by BME/OPT does depend on which social-cost function we use).

However, in contrast to the sc_{wavg} setting (where there is a bound of $\text{BME/OPT} \leq 2$ for n -player games), even mediators cannot enforce outcomes that are close to OPT under sc_{uavg} as the number of players grows, even in linear-latency games. (The construction that follows also demonstrates that none of BCE, BNE, and WNE can provide constant approximations to OPT.)

Lemma 6.2. *Under the social-cost function sc_{uavg} , BME/OPT is not bounded by any constant in 2-machine, n -job weighted load-balancing games with linear latency functions.*

Proof. Consider latency functions $f_L(x) = x$ and $f_R(x) = kx$, and $1 + k^2$ jobs: one “big” job with weight $w_1 = k^3$, and k^2 “small” jobs of unit weight.

If we place the big job on R and the remaining jobs on L , then the total unweighted cost is $2k^4$: there are k^2 players on L with latency k^2 , and one player on R with latency k^4 . Thus $\text{OPT} \leq 2k^4$. (In fact it is not hard to argue that $\text{OPT} = 2k^4$.)

Suppose M is a mediated equilibrium. Let p_R and p_L denote the probabilities that the big job is placed on machine R or L , respectively, under M . The big job incurs a cost of at least k^4 when assigned to machine R , but would incur a cost of at most $k^3 + k^2$ by deterministically playing L . Thus p_R must be $O(1/k)$ to prevent the big job from deviating to L , so p_L must be $1 - o(1)$.

We now argue that when the big job is assigned to machine L , the social cost is large. Indeed, at least $\frac{1}{2}k^2$ of the unit-sized jobs must be assigned to either L or R in any outcome. In the first case, the cost from just those jobs on L is at least $(k^3 + \frac{1}{2}k^2)\frac{1}{2}k^2 \geq \frac{1}{2}k^5$. In the second case, the cost of just those jobs on R is at least $k(\frac{1}{2}k^2)^2 = \frac{1}{4}k^5$. Thus any outcome in which the big job is assigned to L has a cost of $\Omega(k^5)$, and therefore the cost of any mediated equilibrium is $\Omega(k^5)$, and BME/OPT is $\Omega(k)$. \square

7 Mediated Equilibria and the Repeated Cost of Anarchy/Stability

The notion of “punishment” that we see in mediated equilibria is highly reminiscent of the Folk Theorem from the economics literature. The connection between mediated equilibria in one-shot games and Nash equilibria in repeated games turns out to be deep. In this section, we explore this connection, including proofs showing the precise correspondence of BME/OPT in a one-shot game G to the price of stability in the *repeated* game with stage game G . (An analogous equivalence holds between WME/OPT in the one-shot game and the price of anarchy in the repeated game.)

7.1 Mediated Equilibria and Reservation Costs

Let G be an arbitrary n -player game. Each player i selects a strategy from the set S_i ; when the players choose a strategy profile \mathbf{s} , player i pays cost $c_i(\mathbf{s})$. Define the *reservation cost* c_i for player i as the worst

cost that can be imposed on player i by a mixed strategy profile played by the remaining players in the game; formally,

$$c_i = \max_{\sigma_{-i}} \left[\min_{\sigma_i} c_i(\sigma_{-i}, \sigma_i) \right], \quad (7.1)$$

where σ_i denotes a mixed strategy for player i and σ_{-i} denotes a mixed strategy profile for all players except i . Let \underline{c} denote the vector of reservation costs.

Let \mathbf{x} be a vector of costs for the players. The vector \mathbf{x} *Pareto dominates* \underline{c} if $x_i \leq c_i$ for each player i , and *strictly Pareto dominates* \underline{c} if the inequality is strict for every i . We say that \mathbf{x} is *feasible* if it can be expressed as the convex combination of costs of strategy profiles \mathbf{s} —that is, the costs \mathbf{x} can be achieved by an appropriate mixture (possibly correlated) of strategy profiles in G .

Theorem 7.1. *Let \underline{c} be the vector of reservation costs in a game G . There is a mediated equilibrium in G achieving costs \mathbf{x} if and only if \mathbf{x} is feasible and Pareto dominates \underline{c} .*

Proof (\Rightarrow). Suppose that M is a mediated equilibrium in G that achieves the cost vector \mathbf{x} . It is immediate by definition of a mediator that \mathbf{x} is feasible. To show that \mathbf{x} Pareto dominates \underline{c} , consider a particular player i . Let σ_{-i} denote the mixed strategy profile that M uses when all players other than i delegate—that is, let σ_{-i} be the “punishment” profile if player i deviates. Because M is an equilibrium in which player i delegates to receive cost x_i , we know that x_i is less than or equal to her best response to σ_{-i} , which is less than or equal to \underline{c}_i because σ_{-i} is one of the terms in the maximum in (7.1). Thus $x_i \leq \underline{c}_i$.

(\Leftarrow). Let \mathbf{x} be a feasible cost vector that Pareto dominates \underline{c} . Define the mediator M as follows.

- If everyone delegates, then play according to the appropriate mixture of strategy profiles to achieve costs \mathbf{x} . (This mixture exists because \mathbf{x} is feasible.)
- If one player i deviates, then have the remaining players play according to the strategy profile σ_{-i} that achieves the maximum in (7.1) for player i .

(As usual, if more than one player deviates, we have M play arbitrarily.) Player i pays a cost of x_i for delegating, and pays at least \underline{c}_i for deviating, by (7.1). By assumption $x_i \leq \underline{c}_i$, so player i does not wish to deviate from M . Thus M is a mediated equilibrium incurring cost vector \mathbf{x} . \square

Theorem 7.1 implies that the load-balancing analysis in this paper is actually an analysis of the vector of reservation costs \underline{c} . Although we cast our results in different language, all of the work in characterizing BME was based on optimizing the social-cost function over the set of feasible cost vectors \mathbf{x} that Pareto dominate \underline{c} . This rephrasing might appear to trivialize our work—one merely has to find \underline{c} , after all!—but recent work by Borgs et al. [6] shows that computing reservation costs is NP-hard when there are three or more players. Thus, while this perspective on the BME/OPT ratio is valuable, it does not turn the problem into a triviality.

7.2 The Folk Theorem and the Price of Anarchy/Stability in Repeated Games

Let G be an n -player game, with strategy set S_i and cost function c_i for each player i . We refer to G as the *stage game*. Following Fudenberg–Tirole [17], the *repeated game*—more precisely, the *infinitely repeated*

game with discount factor δ —is denoted by $G(\delta)$, for a given value of $\delta \in (0, 1)$. In each stage t , each player i chooses an action from S_i after observing the history of strategy profiles selected in all previous stages. The cost to player i is $(1 - \delta) \cdot \sum_t \delta^t \cdot c_i(\mathbf{s}_t)$, where \mathbf{s}_t denotes the strategy profile selected by the players in stage t . (A player is indifferent between paying a dollar today and paying $1/\delta > 1$ dollars tomorrow; the normalizing constant $1 - \delta$ puts the costs in $G(\delta)$ on the same scale as the costs in G .) A player i can also choose a mixed strategy σ_i over her pure strategies; expected costs in $G(\delta)$ are calculated in the usual way.

The “Folk Theorem” from the game-theory community—so named because it was “folk wisdom” well before it first appeared in print—characterizes the set of Nash equilibria in the repeated game $G(\delta)$. Variants of this theorem appear in a number of different places in the literature (e.g., [4, 16, 33]); an excellent exposition appears in the textbook of Fudenberg and Tirole [17]. (For simplicity, following Fudenberg–Tirole, we assume the presence of a public randomizing device in the repeated game; there are some complications with the convexity of the set of feasible payoffs without additional assumptions.)

Theorem 7.2 (The Folk Theorem (see [17])). *Let $\underline{\mathbf{c}}$ be the vector of reservation costs in a game G , and let \mathbf{x} be a feasible cost vector that strictly Pareto dominates $\underline{\mathbf{c}}$. Then there exists a $\delta^* < 1$ such that, for all $\delta \in (\delta^*, 1)$, there is a Nash equilibrium in $G(\delta)$ that achieves costs \mathbf{x} .*

Intuitively, the Nash equilibrium is structured as follows. As long as everyone plays according to the strategy profile that produces costs \mathbf{x} , then everyone achieves the desired costs. If a player i deviates in stage t , then all other players play according to the maximally punishing strategy profile given by (7.1) in round $t + 1$ and forever after. For sufficiently patient players ($\delta \approx 1$), the benefit in round t is swamped by the perpetual losses in round $t + 1$ and beyond.

The only real difference between Theorems 7.1 and 7.2 is on the boundary, when \mathbf{x} Pareto dominates $\underline{\mathbf{c}}$, but does not *strictly* Pareto dominate $\underline{\mathbf{c}}$. The following lemma summarizes the difference:

Lemma 7.3. *Let G be an n -player game with reservation costs $\underline{\mathbf{c}}$. Define the sets of cost vectors $X^\bullet = \{\text{feasible } \mathbf{x} : \mathbf{x} \text{ Pareto dominates } \underline{\mathbf{c}}\}$ and $X^\circ = \{\text{feasible } \mathbf{x} : \mathbf{x} \text{ strictly Pareto dominates } \underline{\mathbf{c}}\}$. Then:*

1. *if $\mathbf{x} \in X^\circ$, then the costs \mathbf{x} are achieved by some mediated equilibrium in G and by some Nash equilibrium in $G(\delta)$, for some $\delta < 1$.*
2. *if $\mathbf{x} \in X^\bullet - X^\circ$, then the costs \mathbf{x} are achieved by some mediated equilibrium in G .*
3. *if $\mathbf{x} \notin X^\bullet$, then the costs \mathbf{x} are not achieved by any mediated equilibrium in G nor by any Nash equilibrium in $G(\delta)$, for any $\delta < 1$.*

Proof. The first two claims follow immediately from Theorem 7.1 and the Folk Theorem. For the third claim, it is easy to show that if \mathbf{x} is not feasible in the stage game G , then there is no Nash equilibrium achieving costs \mathbf{x} in the repeated game $G(\delta)$, for any δ . Furthermore, if there exists a player i for which $x_i > \underline{c}_i$, then there is no Nash equilibrium achieving costs \mathbf{x} in the repeated game $G(\delta)$, for any δ : player i can play a best response in every stage, thereby achieving cost at most \underline{c}_i . (See Fudenberg–Tirole [17] for a detailed proof.) \square

We can now use this result to relate the price of anarchy/stability in the repeated game $G(\delta)$ to the ratio of the worst/best mediated equilibrium to OPT in G . The only case in which the correspondence is

not necessarily exact is when X° is empty; otherwise, the difference in boundary conditions is irrelevant because every \mathbf{x} achievable by a mediated equilibrium in G is arbitrarily close to a Nash equilibrium in $G(\delta)$, for $\delta \approx 1$. We say that a game G is *nondegenerate* if X^\bullet has nonzero volume, or, equivalently, if X° is nonempty.

Theorem 7.4. *Suppose that G is nondegenerate. As δ tends to 1, the price of anarchy of $G(\delta)$ tends to $\text{WME}(G)/\text{OPT}(G)$ and the price of stability of $G(\delta)$ tends to $\text{BME}(G)/\text{OPT}(G)$.*

Proof. By Lemma 7.3, for all δ the Nash equilibria of $G(\delta)$ are a subset of X^\bullet , which is the set of outcomes achievable as mediated equilibria. Because X^\bullet is compact, $\text{WME}(G)$ is achieved somewhere on X^\bullet . Hence, for all δ , $\text{WNE}(G(\delta)) \leq \text{WME}(G)$. Now let $\varepsilon > 0$. Because the social cost function is continuous on X^\bullet , there exists an $\mathbf{x} \in X^\circ$ such that $\text{WME}(G) - \mathbf{x} < \varepsilon$. Let $\delta^* = \delta^*(\mathbf{x})$ be given by Theorem 7.2. Then for all $\delta \in (\delta^*, 1)$,

$$\text{WNE}(G(\delta)) \geq \mathbf{x} > \text{WME}(G) - \varepsilon.$$

Thus

$$\lim_{\delta \rightarrow 1^-} \text{WNE}(G(\delta)) = \text{WME}(G),$$

which implies the price of anarchy claim. The price of stability argument is similar. \square

Thus our analysis of the quality of mediated equilibria in load-balancing games also directly implies the same results on the quality of equilibria in repeated load-balancing games.

8 Future Directions

In this paper we have begun to analyze the power of mediators in the spirit of the price of stability, focusing on load-balancing games under the weighted average latency social-cost function. We have a complete story for unweighted games and for weighted games with general latency functions. The biggest open question is the gap between BME and OPT in n -player weighted linear games. We know that for all such games $\text{BME}/\text{OPT} \leq 2$, and that there exist examples in which $\text{BME}/\text{OPT} \approx 1.2071$. What is the worst-case BME/OPT for $n \geq 3$ players? In this context, it would also be interesting to understand the depth of the connection between the weighted and unweighted total latency social-cost functions: the fact that the same instance is the worst case for both functions in the 2-player case (from Theorem 4.1 and Theorem 6.1) was unexpected and may suggest a fruitful line for future work. The relationship to the Folk Theorem may also be helpful in this analysis; we need only understand reservation costs and feasible points to improve these bounds.

The broader direction for future research, of course, is to characterize the power of mediators in games beyond load balancing. It is an interesting question as to how much better mediated equilibria are than correlated equilibria in, say, linear-latency weighted congestion games.

References

- [1] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 279–287, 2007. [3, 5](#)

- [2] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. On the value of correlation. *Journal of Artificial Intelligence*, 33:575–613, 2008. [5](#)
- [3] Robert J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987. [2](#)
- [4] Robert J. Aumann and Lloyd S. Shapley. Long term competition: A game theoretic analysis. UCLA Economics Working Papers 676, UCLA Department of Economics, 1992. [21](#)
- [5] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005. [4](#), [5](#), [10](#)
- [6] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Adam Tauman Kalai, Vahab Mirrokni, and Christos Papadimitriou. The myth of the folk theorem. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*, pages 365–372, 2008. [5](#), [20](#)
- [7] Ioannis Caragiannis. Personal communication, March 2010. [4](#), [5](#), [10](#), [16](#)
- [8] Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. Tight bounds for selfish and greedy load balancing. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 311–322, 2006. [4](#), [5](#), [8](#), [10](#), [14](#)
- [9] George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the Annual European Symposium on Algorithms (ESA)*, pages 59–70, 2005. [4](#), [5](#)
- [10] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005. [4](#), [5](#), [10](#)
- [11] Artur Czumaj and Berthold Vöcking. Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms*, 3(1):4, 2007. [5](#)
- [12] Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. Convergence time to Nash equilibria. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 502–513, 2003. [6](#), [17](#)
- [13] Chaim Fershtman, Kenneth L. Judd, and Ehud Kalai. Observable contracts: Strategic delegation and cooperation. *International Economic Review*, 32(3):551–559, August 1991. [5](#)
- [14] Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2):226–239, 2005. [4](#), [5](#), [16](#)

- [15] Dimitris Fotakis, Spyros C. Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 123–134, 2002. [6](#), [17](#)
- [16] D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 1986. [21](#)
- [17] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1993. [4](#), [20](#), [21](#)
- [18] M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash equilibria, the price of anarchy and the fully mixed Nash equilibrium conjecture. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 51–65, 2005. [5](#)
- [19] Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. Nash equilibria in discrete routing games with convex latency functions. *Journal of Computer and System Sciences*, 74(7):1199–1225, 2008. [5](#)
- [20] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000. [5](#)
- [21] Sham Kakade, Michael Kearns, John Langford, and Luis Ortiz. Correlated equilibria in graphical games. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 42–47, 2003. [5](#)
- [22] Adam Tauman Kalai, Ehud Kalai, Ehud Lehrer, and Dov Samet. A commitment folk theorem. *Games and Economic Behavior*, 69(1):127–137, 2010. [5](#)
- [23] Anshul Kothari, Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Congestion games, load balancing, and price of anarchy. In *Proceedings of the Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN)*, 2004. [4](#), [5](#)
- [24] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999. [2](#), [4](#), [5](#)
- [25] Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. A new model for selfish routing. *Theoretical Computer Science*, 406(3):187–206, 2008. [4](#), [5](#)
- [26] Marios Mavronicolas and Paul Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007. [5](#)
- [27] Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009. [3](#), [4](#), [5](#), [6](#), [7](#)
- [28] Christos H. Papadimitriou and Tim Roughgarden. Computing equilibria in multi-player games. *Journal of the ACM*, 55(3):1–29, 2008. [5](#)

- [29] Bezalel Peleg and Ariel D. Procaccia. Implementation by mediated equilibrium. Discussion paper 463, Center for the Study of Rationality, Hebrew University, September 2007. 3, 5
- [30] Bezalel Peleg and Ariel D. Procaccia. Mediators and truthful voting. Working paper, June 2008. 3, 5
- [31] Ola Rozenfeld and Moshe Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *Proceedings of the Workshop on Internet and Network Economics (WINE)*, pages 74–86, 2006. 3, 5
- [32] Ola Rozenfeld and Moshe Tennenholtz. Routing mediators. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1488–1493, 2007. 3, 5, 7
- [33] A. Rubenstein. Finite automata play the repeated prisoners dilemma. *Journal of Economic Theory*, 39:83–96, 1986. 21
- [34] Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007. 4, 5
- [35] Moshe Tennenholtz. Game-theoretic recommendations: Some progress in an uphill battle. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, pages 10–16, 2008. 3, 5
- [36] Berthold Vöcking. Selfish load balancing. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007. 5, 6, 17

A Omitted Proofs: Unweighted Average Social Cost

Lemma A.1. *The social optimum with respect to the social-cost function sc_{uavg} has cost*

$$\text{OPT} = \begin{cases} 2(w+1) & \text{if } a \geq w+2 \\ a+w & \text{otherwise.} \end{cases}$$

Proof. We claim that the social optimum is never uniquely achieved by *RR* or *LR*. The social cost of *RR* is $2a(w+1)$, which is $a \geq 1$ times the social cost of *LL*. Similarly, the social cost of *RL* is $a+w = a+w-1+1 \leq a+a(w-1)+1 = aw+1$, the social cost of *LR*. Thus the social optimum is achieved by either *LL* or *RL*, at total social cost $2(1+w)$ or $a+w$. Observe that $2(w+1) = 2w+2$ exceeds $a+w$ if and only if $a < w+2$. \square

Lemma A.2. *The best Nash equilibrium and the best correlated equilibrium with respect to the social-cost function sc_{uavg} have cost*

$$\text{BNE} = \text{BCE} = \begin{cases} 2(w+1) & \text{if } a > w+1 \\ a+w & \text{otherwise.} \end{cases}$$

Proof. Suppose $a > w + 1$. Then L is a dominant strategy for both players. Thus the only Nash equilibrium is for both to play L . Furthermore, no correlated equilibrium can play any outcome other than LL with nonzero probability, because R is a dominated strategy.

On the other hand, suppose $a \leq w + 1$. Then RL is a Nash equilibrium: the little player prefers the cost of a to LL 's cost of $w + 1$; and the big player certainly prefers w to $a(w + 1)$. There cannot be a better Nash equilibrium or correlated equilibrium for $a \leq w + 1 < w + 2$ by Theorem A.1. \square

Lemma A.3. *Suppose $a > w + 1$. In the best mediated equilibrium, if one player deviates from the mediator then the mediator will have the one remaining delegating player choose L . Under this mediator, the cost incurred by a deviating player is $1 + w$.*

Proof. The proof is precisely that of Theorem 4.4, as there is no dependence on the social-cost function. \square

Lemma A.4. *Suppose $a > w + 1$. The mediator $M = \langle \langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle \rangle$ forms a mediated equilibrium if and only if*

$$w + 1 \geq (w + 1)p_{LL} + p_{LR} + ap_{RL} + a(w + 1)p_{RR} \quad (\text{A.1})$$

$$w + 1 \geq (w + 1)p_{LL} + awp_{LR} + wp_{RL} + a(w + 1)p_{RR} \quad (\text{A.2})$$

and the social cost of M with respect to the social-cost function sc_{avg} is the (unweighted) sum of the right-hand sides of (A.1) and (A.2).

Proof. By the definition of the mediator and by Theorem A.3, the cost to a nondelegating player is $w + 1$. The costs to the two players if both delegate are the right-hand sides of (A.1) and (A.2), respectively. Thus M is a mediated equilibrium if and only if both constraints are satisfied. \square

Lemma A.5. *Suppose $a > w + 1$. If $M = \langle \langle p_{LL}, p_{LR}, p_{RL}, p_{RR} \rangle \rangle$ is a mediated equilibrium, then so is the mediator $M' = \langle \langle p_{LL} + p_{RR}, p_{LR}, p_{RL}, 0 \rangle \rangle$ that shifts all the probability of playing RR to the outcome LL . Furthermore, the social cost of the M' with respect to the social-cost function sc_{avg} is no larger than the social cost of M .*

Proof. Suppose that M is a mediated equilibrium—in other words, by Theorem A.4, suppose that the probabilities p_{LL} , p_{LR} , p_{RL} , and p_{RR} satisfy (A.1) and (A.2). Because $a(w + 1) \geq w + 1$, the right-hand sides of both (A.1) and (A.2) have not increased when we move from M to M' , so the social cost has not increased while the constraints remain satisfied. Thus M' has social cost that is no greater than the social cost of M and remains a mediated equilibrium. \square

Lemma A.6. *Suppose $a > w + 1$. For any two probabilities $p > 0$ and $q \geq 0$, define the mediator $M_{p,q} = \langle \langle 1 - p, pq, p(1 - q), 0 \rangle \rangle$.*

- (i) *If $a > w + 1 + \frac{1}{w}$, then no $M_{p,q}$ is a mediated equilibrium.*
- (ii) *If $a \leq w + 1 + \frac{1}{w}$, then there are values of p, q for which $M_{p,q}$ is a mediated equilibrium. Furthermore, the mediated equilibrium $M_{p,q}$ with the lowest social cost is achieved when $q = (a - w - 1)/(a - 1)$ and $p = 1$, when the cost is $aw - w^2 + w + 1$.*

Proof. Applying Theorem A.4 to the mediator $M_{p,q} = \langle \langle 1-p, pq, p(1-q), 0 \rangle \rangle$ implies that $M_{p,q}$ is a mediated equilibrium if and only if

$$\begin{aligned} w+1 &\geq (w+1)(1-p) + ap(1-q) + pq \\ w+1 &\geq (w+1)(1-p) + wp(1-q) + awpq. \end{aligned}$$

Collecting like terms and dividing both inequalities by p , we have

$$w+1 \geq a(1-q) + q \qquad w+1 \geq w(1-q) + awq.$$

Solving for q yields

$$\frac{a-w-1}{a-1} \leq q \leq \frac{1}{(a-1)w}. \quad (\text{A.3})$$

Note that there is a q satisfying (A.3) if and only if

$$\frac{a-w-1}{a-1} \leq \frac{1}{(a-1)w} \iff a \leq w+1 + \frac{1}{w}.$$

Thus if $a > w+1 + \frac{1}{w}$, then no $M_{p,q}$ is a mediated equilibrium. For the case in which $a \leq w+1 + \frac{1}{w}$, note that the social cost of $M_{p,q}$ is given by

$$\begin{aligned} \text{social cost of } M_{p,q} &= (1-p)2(w+1) + pq(1+aw) + p(1-q)(a+w) \\ &= 2(w+1) + p(a-w-2) + (a-1)(w-1)pq. \end{aligned} \quad (\text{A.4})$$

By assumption, we have $a \geq 1$, $w \geq 1$, and $p > 0$, so this social cost is an increasing function of q . Thus the social cost is minimized for the smallest value of q such that (A.3) is satisfied, namely $q^* = (a-w-1)/(a-1)$. Plugging this value for q^* into (A.4), we have that

$$\begin{aligned} \text{social cost of } M_{p,q^*} &= 2(w+1) + p(a-w-2) + (w-1)p(a-w-1) \\ &= 2w+2 + p(aw-w^2-w-1). \end{aligned} \quad (\text{A.5})$$

By assumption, we have that $a \leq w+1 + \frac{1}{w}$, and thus

$$aw-w^2-w-1 \leq (w+1 + \frac{1}{w})w-w^2-w-1 = w^2+w+1-w^2-w-1 = 0.$$

Because $w > 0$, the social cost in (A.5) is a decreasing function of p , so the social cost is minimized for the largest value of p , namely $p = 1$. In this case, the social cost is

$$2w+2 + aw-w^2-w-1 = aw-w^2+w+1$$

and we are done. □

Lemma A.7. *The best mediated equilibrium has cost*

$$\text{BME} = \begin{cases} 2(w+1) & \text{if } a > w+1 + \frac{1}{w} \\ aw-w^2+w+1 & \text{if } a \in (w+1, w+1 + \frac{1}{w}] \\ a+w & \text{if } a \leq w+1. \end{cases}$$

Proof. By Theorem A.3, Theorem A.5, and Theorem A.6, there are only two candidates for the best mediated equilibrium: one of the form $\langle\langle 0, q, 1 - q, 0 \rangle\rangle$ as in Theorem A.6, or one of the form $\langle\langle 1, 0, 0, 0 \rangle\rangle$, which is the only mediator that does not meet the hypotheses of Theorem A.6 after the transformations of Theorem A.3 and Theorem A.5.

For $a > w + 1 + \frac{1}{w}$, by Theorem A.6 there is no mediated equilibrium of the form $\langle\langle 0, q, 1 - q, 0 \rangle\rangle$, so the only mediated equilibrium is the $\langle\langle 1, 0, 0, 0 \rangle\rangle$ mediator, which is in fact the pure Nash equilibrium from Theorem A.2, with social cost $2(w + 1)$.

For $w + 1 < a \leq w + 1 + \frac{1}{w}$, the best $\langle\langle 0, q, 1 - q, 0 \rangle\rangle$ mediated equilibrium has social cost $aw - w^2 + w + 1$ by Theorem A.6. For values of a in this range, we have that $aw - w^2 + w + 1 \leq 2(w + 1)$, by the same argument as in the proof of Theorem A.6, so this mediator achieves a better social cost than the $\langle\langle 1, 0, 0, 0 \rangle\rangle$ mediator.

For $a \leq w + 1$, the best Nash equilibrium achieves the social optimum, by Theorem A.1 and Theorem A.2. Any Nash equilibrium is a mediated equilibrium, and no mediated equilibrium can outperform the social optimum; thus we are done. \square

The results, in summary, are the following:

	$a \in [1, w + 1]$	$a \in (w + 1, w + 1 + \frac{1}{w}]$	$a \in (w + 1 + \frac{1}{w}, 2 + w]$	$a \in (2 + w, \infty)$
OPT	$a + w$	$a + w$	$a + w$	$2w + 2$
BME	$a + w$	$aw - w^2 + w + 1$	$2w + 2$	$2w + 2$
BCE = BNE	$a + w$	$2w + 2$	$2w + 2$	$2w + 2$

Lemma A.8. *The ratio between the cost of the best mediated equilibrium and the social optimum is at most $\frac{2w^2 + 2w}{2w^2 + w + 1}$, and that bound is tight when $a = w + 1 + \frac{1}{w}$.*

Proof. By Theorem A.1 and Theorem A.7, we have the cost of the social optimum and of the best mediated equilibrium. We will consider the various ranges for a defined by those lemmas.

If $a \leq w + 1$, then both the social optimum and the best mediated equilibrium have social cost $a + w$, for a ratio of one.

If $w + 1 < a \leq w + 1 + \frac{1}{w}$, then the best mediated equilibrium has cost $aw - w^2 + w + 1$ and the social optimum has cost $a + w$. Thus the ratio is

$$\frac{aw - w^2 + w + 1}{a + w} = \frac{aw + w^2 - 2w^2 + w + 1}{a + w} = w + \frac{-2w^2 + w + 1}{a + w}.$$

Because $w \geq 1$, we have that $-2w^2 + w + 1 \leq 0$, and thus this ratio is an increasing function of a . Therefore the ratio throughout this range is upper-bounded by its value at $a = w + 1 + \frac{1}{w}$, when its value is

$$\frac{(w + 1 + \frac{1}{w})w - w^2 + w + 1}{w + 1 + \frac{1}{w} + w} = \frac{2w + 2}{2w + 1 + \frac{1}{w}} = \frac{2w^2 + 2w}{2w^2 + w + 1}. \quad (\text{A.6})$$

If $w + 1 + \frac{1}{w} < a < w + 2$, then the best mediated equilibrium has cost $2w + 2$ and the social optimum has cost $a + w$. Thus the ratio of their costs is $\frac{2w + 2}{a + w}$, which is a decreasing function of a . Therefore the ratio throughout this range is upper-bounded by its value at $a = w + 1 + \frac{1}{w}$. For this value of a , we have

$aw - w^2 + w + 1 = 2w + 2$, and thus the ratio is upper-bounded throughout this range by $\frac{2w^2+2w}{2w^2+w+1}$, just as in (A.6).

If $a \geq w + 2$, then both the social optimum and the best mediated equilibrium have social cost $2(w + 1)$, for a ratio of one. \square

Theorem A.9. *Under unweighted total latency social cost, in all linear weighted 2-machine 2-job load-balancing games:*

- $\text{BME}/\text{OPT} \leq \frac{2+4\sqrt{2}}{7} \approx 1.0938$. This bound is tight for jobs of weights 1 and $1 + \sqrt{2}$ and for latency functions $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x$.
- $\text{BCE}/\text{BME} \leq \frac{4}{3}$. The bound is tight for the unweighted 2-job, 2-machine example with $f_L(x) = x$ and $f_R(x) = (2 + \varepsilon) \cdot x$ from Section 3.

Proof. Without loss of generality, we scale the jobs and latency functions so that the jobs' weights are 1 and $w \geq 1$ and the machines' latency functions are $f_L(x) = x$ and $f_R(x) = ax$ for $a \geq 1$. By Theorem A.8, the worst-case ratio of the best mediated equilibrium to the social optimal outcome occurs when $a = w + 1 + 1/w$, when

$$\frac{\text{BME}}{\text{OPT}} = \frac{2w^2 + 2w}{2w^2 + w + 1}.$$

We wish to upper-bound this ratio over all $w \geq 1$. By simple calculus, the ratio is maximized when $w^2 - 2w - 1 = 0$, or when $w = 1 + \sqrt{2}$. For this value of w , we have $\text{BME}/\text{OPT} = \frac{8+6\sqrt{2}}{8+5\sqrt{2}} = \frac{2+4\sqrt{2}}{7}$. Thus the ratio of the best mediated equilibrium to OPT optimum never exceeds $\frac{2+4\sqrt{2}}{7}$.

Furthermore, for the load-balancing game with jobs of weights 1 and $1 + \sqrt{2}$, and with machines with latency functions $f_L(x) = x$ and $f_R(x) = (1 + 2\sqrt{2}) \cdot x \approx 3.8284x$, the ratio of the best mediated equilibrium to the social optimum is $\frac{2+4\sqrt{2}}{7} \approx 1.0938$.

The worst-case value of a for BCE/BME is again $a = 1 + w + \varepsilon$, when $\text{BCE}/\text{BME} = \frac{2w+2}{2w+1+\varepsilon}$, which is again maximized at $w = 1$ when the ratio is $4/3$. \square

AUTHORS

Joshua R. Davis
 Department of Mathematics
 Carleton College
 Northfield, MN 55057
 jd@joshuadavis.us
<http://www.joshuadavis.us/>

David Liben-Nowell
Department of Computer Science
Carleton College
Northfield, MN 55057
dlibenno@carleton.edu
<http://cs.carleton.edu/faculty/dlibenno/>

Alexa Sharp
Department of Computer Science
Oberlin College
Oberlin, OH 44074
alexa.sharp@oberlin.edu
<http://www.cs.oberlin.edu/~asharp/>

Tom Wexler
Department of Computer Science
Oberlin College
Oberlin, OH 44074
tom.wexler@oberlin.edu
<http://www.cs.oberlin.edu/~wexler/>