Restoring Pure Equilibria to Weighted Congestion Games

KONSTANTINOS KOLLIAS, Stanford University
TIM ROUGHGARDEN, Stanford University

Congestion games model several interesting applications, including routing and network formation games, and also possess attractive theoretical properties, including the existence of and convergence of natural dynamics to a pure Nash equilibrium. Weighted variants of congestion games that rely on sharing costs proportional to players' weights do not generally have pure-strategy Nash equilibria. We propose a new way of assigning costs to players with weights in congestion games that recovers the important properties of the unweighted model. This method is derived from the Shapley value, and it always induces a game with a (weighted) potential function. For the special cases of weighted network cost-sharing and weighted routing games with Shapley value-based cost shares, we prove tight bounds on the worst-case inefficiency of equilibria. For weighted network cost-sharing games we precisely calculate the price of stability for any given player weight vector, while for weighted routing games we precisely calculate the price of anarchy, as a parameter of the set of allowable cost functions.

Categories and Subject Descriptors: G.2.1 [**Discrete Mathematics**]: Combinatorics; G.2.2 [**Graph Theory**]: Network problems

General Terms: Economics, Theory

Additional Key Words and Phrases: Weighted congestion games, weighted Shapley value, price of anarchy, price of stability

ACM Reference Format:
Kollias, K., and Roughgarden, T. 2014. Restoring Pure Equilibria to Weighted Congestion Games. ACM 0, 0, Article 0 (2014), 24 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION.

Congestion games are a well-studied generalization of several game-theoretic models, including some fundamental network formation games and routing games. In the standard model [Rosenthal 1973a], there is a ground set of resources, and each player has a set of allowable strategies, each of which is a subset of resources. For example, the strategies of a player could correspond to the paths of a network with a given source and sink. Each resource has a per-user cost that depends on the number of players that use it, and the goal of each player is to minimize the sum of the resources' costs in its strategy, given the strategies chosen by the other players. In atomic selfish routing games [Rosenthal 1973b; Roughgarden and Tardos 2002], strategies correspond to paths and the per-unit cost function $c_e(\cdot)$ of each resource e is nondecreasing. In network cost-sharing games [Anshelevich et al. 2008], strategies correspond to paths

A preliminary version of this work appeared in the ICALP '11 conference proceedings.
Author's addresses: K. Kollias. Management Science and Engineering Department, Stanford University. T. Roughgarden. Computer Science Department, Stanford University.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

and the (decreasing) cost functions have the form $c_e(x_e) = \gamma_e/x_e$, where γ_e is the fixed installation cost of edge e and x_e is the number of players that share it.

A *pure Nash equilibrium (PNE)* is a strategy profile such that no player can decrease its cost via a unilateral deviation. Many games, such as "Rock-Paper-Scissors", have no PNE. Rosenthal [Rosenthal 1973a] used a potential function argument to show that every congestion game — and thus every atomic selfish routing and network cost-sharing game — has at least one PNE. Moreover, better-response dynamics is guaranteed to converge to a PNE [Monderer and Shapley 1996].

Every player of a congestion game imposes the same load on a resource. There are many motivations for relaxing this assumption and allowing non-uniform resource consumption. For example, in a network context, players could have different durations of resource usage, different bandwidth requirements, or different contracts with the network provider. Almost all research to date has modeled non-uniform players in congestion-type games through *proportional cost sharing* [Aland et al. 2011; Anshelevich et al. 2008; Awerbuch et al. 2005; Bhawalkar et al. 2010; Chen and Roughgarden 2009; Gairing and Schoppmann 2007; Harks and Klimm 2012; Harks et al. 2011; Milchtaich 1996; Monderer and Shapley 1996]. The first assumption in proportional cost sharing is that each player i has a positive weight w_i, with larger weights indicating larger resource usage. To explain the second assumption in a general way, let $C_e(S_e)$ denote the joint cost incurred by the subset S_e of users of the resource e. For example, in a network cost-sharing game, $C_e(S_e)$ is the fixed cost γ_e provided S_e is non-empty (and is 0 otherwise). In (weighted) atomic selfish routing, $C_e(S_e)$ is $x_e \cdot c_e(x_e)$, where $c_e(\cdot)$ is the per-flow unit resource cost function and $x_e = \sum_{i \in S_e} w_i$ is the total weight of the players using e. Proportional cost sharing dictates that each player $i \in S_e$ pays a $w_i / \sum_{j \in S_e} w_j$ fraction of $C_e(S_e)$ for the resource e.

Unfortunately, most of the attractive theoretical properties of congestion games do not carry over to their weighted counterparts with proportional cost sharing. Network cost-sharing games with at least three players need not have a PNE [Chen and Roughgarden 2009]. Even when PNE do exist in such games, they can be much costlier (relative to an optimal solution) than in the unweighted case [Anshelevich et al. 2008; Chen and Roughgarden 2009]. Atomic selfish routing games with weighted players do not generally have PNE [Goemans et al. 2005; Harks and Klimm 2012; Rosenthal 1973b], except when all cost functions are affine [Fotakis et al. 2005] and in some other isolated special cases [Harks and Klimm 2012].

Guaranteed existence of PNE is an important property. There are, of course, more general equilibrium concepts like the mixed-strategy Nash equilibrium that are guaranteed to exist in every finite game, but randomized solution concepts suffer from well-known drawbacks in interpretation and implementation (see e.g. [Osborne and Rubinstein 1994, §3.2]). Particularly when designing or influencing the game being played, there is good reason to make design decisions that guarantee the existence of and convergence of natural dynamics to a PNE. Previous works have studied how to design systems with such guarantees in the domains of queuing [Mosk-Aoyama and Roughgarden 2009; Moulin 2008; Shenker 1995], network cost-sharing [Chen et al. 2010; Gopalakrishnan et al. 2013], and distributed resource allocation [Marden and Wierman 2013].

1.1. Our Contributions.

We propose a new way of assigning costs to players with weights in congestion-type games, which is derived from the Shapley value. We call the resulting class of games *SV weighted congestion games*. Extending work of Hart and Mas-Colell [Hart and Mas-Colell 1989], we show that every SV weighted congestion game admits a (weighted)

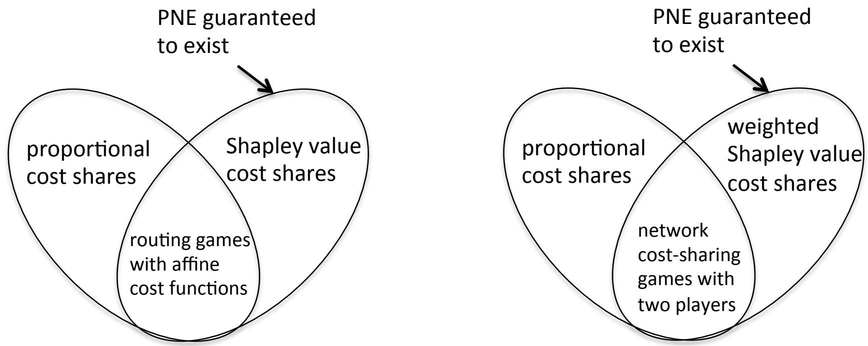

(a) Atomic Routing Games

(b) Network Cost-Sharing Games

Fig. 1. Comparison of traditional proportional cost shares with the Shapley value cost shares proposed in the present work.

potential function. The existence of and convergence of natural dynamics to a PNE in every such game follow immediately.

For example, for the special case of atomic selfish routing games, we derive the cost shares for the users S_e of edge e by applying the standard Shapley value (defined in the next section) to the cost function $C_e(\cdot)$ above with the player set S_e. Since the incremental effect of a player on the joint cost is increasing in its weight, so is its cost share. These Shapley value-based cost shares coincide with proportional shares when all per-user cost functions are affine, but not otherwise (Figure 1(a)). These results explain the previously mysterious fact that the traditional proportional cost shares always yield a potential game when cost functions are affine [Fotakis et al. 2005; Harks and Klimm 2012].

For the special case of network cost-sharing games, the symmetric joint cost function $C_e(\cdot)$ is insensitive to players' weights. To introduce weight-dependent cost shares, we use the *weighted* Shapley value [Kalai and Samet 1987; Shapley 1953], which averages over orderings of the players in a non-uniform way (see the next section for a definition). The resulting cost shares are increasing in weight, and coincide with proportional shares (for all weight vectors) if and only if there are at most two players (Figure 1(b)). These facts explain why, with proportional cost shares, PNE always exist with two players [Anshelevich et al. 2008] but not with at least three [Chen and Roughgarden 2009].

We also provide tight bounds on the inefficiency of equilibria in SV weighted network cost-sharing and atomic selfish routing games. For weighted atomic selfish routing games, we give tight bounds on the worst-case price of anarchy (POA) [Koutsoupias and Papadimitriou 1999] — the ratio between the cost of the worst PNE and of an optimal outcome — with respect to every set of convex cost functions and a worst-case set of player weights. This worst-case POA is slightly larger than that in weighted congestion games with proportional cost shares that have PNE. For example, in routing games with cost functions that are polynomials with degree at most d and nonnegative coefficients, the POA with proportional cost shares is $\approx (c_1 \cdot d/\ln d)^{d+1}$ (when PNE exist) [Aland et al. 2011] and with Shapley value cost shares is $\approx (c_2 \cdot d)^{d+1}$, where $c_1 \approx 1.3$ and $c_2 \approx 0.9$ are constants independent of d. See also Table I. We establish these POA upper bounds with a "smoothness proof" in the sense of Roughgarden [Roughgarden 2009], so these upper bounds apply more generally to all mixed Nash, correlated, and

Table I. The POA in weighted routing games with polynomial cost functions with nonnegative coefficients, for proportional cost shares (when PNE exist) and for Shapley cost shares.

Degree	Proportional	Shapley
1	2.618	2.618
2	9.909	12.626
3	47.82	101.58
4	277.0	$1,117.78$
5	$1,858$	$15,195$
6	$14,099$	$244,399$
7	$118,926$	$4,536,010$
8	$1,101,126$	$95,410,300$
d	$\Theta(\frac{d}{\log d})^{d+1}$	$\Theta(d)^{d+1}$

coarse correlated equilibria of these games. Thus, Shapley cost shares restore PNE to weighted routing games at the expense of modestly increasing inefficiency.

For network cost-sharing games, we focus on the price of stability (POS) [Anshelevich et al. 2008], which is the ratio between the cost of the best PNE and of an optimal solution. The worst-case POA is uninteresting in such games because it equals k, the number of players, no matter how players' cost shares are defined [Chen et al. 2010, Proposition 4.12]. Our main result here is a characterization of the POS as a function of the weight vector w. For every w, we give an explicit lower bound on the POS and prove a matching upper bound for all networks. The special case of $w = (1, 1, \ldots, 1)$ — where the worst-case POS is the kth Harmonic number — is one of the main results in Anshelevich et al. [Anshelevich et al. 2008]. Our lower bound is a straightforward extension of that in [Anshelevich et al. 2008], but our matching upper bound requires a fundamentally new argument. The upper bound in [Anshelevich et al. 2008] for unweighted players follows directly from the proximity between the potential and objective functions; with weighted players, the difference between these two functions can be arbitrarily larger than the POS. Our characterization implies, for example, that the POS remains $O(\log k)$ if players' weights differ by a constant factor, and is $O(\sqrt{k})$ when $w_i = i$ for $i = 1, 2, \ldots, k$. With proportional cost shares, when PNE exist, the POS can be as large as the sum of the players' weights (assuming that $\min_i w_i = 1$) [Chen and Roughgarden 2009]. In this sense, weighted Shapley cost shares both restore PNE to weighted network cost-sharing games and decrease the inefficiency of such equilibria.

2. THE WEIGHTED SHAPLEY VALUE.

We first recall the weighted Shapley value [Kalai and Samet 1987; Shapley 1953]. Consider a set S of players and a cost function $C : 2^S \to \mathbb{R}$. (For us, S is the users of a resource and $C(T)$ is the joint cost that would be incurred if the players of $T \subseteq S$ were its sole users.) For a given ordering π of the players, let $\Delta_i(\pi)$ denote $C(S^i(\pi) \cup \{i\}) - C(S^i(\pi))$, where $S^i(\pi)$ denotes the players preceding i in π.

Each player has a positive weight w_i and a sampling parameter λ_i set to $1/w_i$ [Kalai and Samet 1987; Shapley 1953]. We use the λ_i's to define a distribution over orderings of players, as follows. (When all λ_i's are equal, we recover the uniform distribution and the usual Shapley value.) We first choose the final player in the ordering, with probabilities proportional to the λ_i's; given this choice, we choose the penultimate player randomly from the remaining ones, again with probabilities proportional to the λ_i's; and so on. The weighted Shapley value of player i is defined as the expected value of $\Delta_i(\pi)$ with respect to this distribution over orderings π.

An equivalent way of defining this distribution over orderings of the players is the following. With each player i, we associate a random variable $\mathbf{X}_i$ that is exponentially distributed with rate λ_i. The realized ordering of the players is then considered to be the one that has the corresponding $\mathbf{X}_i$ variables in decreasing order (i.e., the player with the largest $\mathbf{X}_i$ value is the first in the ordering and the player with the smallest $\mathbf{X}_i$ value is the last in the ordering). The intuition behind the equivalence of the two definitions is the following. Suppose we observe the values of the $\mathbf{X}_i$ variables as they increase from 0, until they reach their realized values. The probability that some given $\mathbf{X}_i$ is the next to stop increasing is proportional to the rate of the corresponding player, i.e., $\lambda_i / \sum_{j \in S} \lambda_j$, due to the fact that the random variables are exponential. (See [Kalai and Samet 1987] for more details and a more formal treatment.)

3. CONGESTION GAMES WITH SHAPLEY VALUE COST SHARES.

Sections 3.1 and 3.2 propose novel cost shares with weighted players in network cost-sharing games and routing games, respectively, which ensure the existence of pure-strategy Nash equilibria. Section 3.3 explains the general construction for arbitrary congestion games.

3.1. Network Cost-Sharing Games.

In an *SV network cost-sharing game*, each player $i = 1, 2, \ldots, k$ has a weight $w_i \geq 1$ and a sampling parameter $\lambda_i = 1/w_i$. We can assume that $w_1 \leq w_2 \leq \ldots \leq w_k$ and we do so for the rest of the paper. Player i aims to construct a path P_i from a given node s_i to a given node t_i in a directed graph $G = (V, E)$, where every $e \in E$ has a fixed nonnegative cost γ_e. With respect to a fixed path vector P, we write S_e for the users of edge e. The cost function C_e corresponding to edge e is $C_e(S_e) = \gamma_e$ if $S_e \neq \emptyset$ and $C_e(S_e) = 0$ if $S_e = \emptyset$.

We next give a probabilistic representation of weighted Shapley cost shares and the corresponding potential function, in terms of independent exponentially distributed random variables, as in the last paragraph of Section 2. Let T be a subset of the players. For every player $i \in T$, let $\mathbf{X}_i$ be an exponentially distributed random variable with rate λ_i. We then define the per-unit weighted Shapley share of i on an edge e used by the players T as the probability that X_i is the largest random variable among those associated with T.

Definition 3.1. In an SV network cost-sharing game, the weighted Shapley share of player $i \in S_e$ for using the edge e is

$$\chi_{i,e}(S_e) = \gamma_e \cdot \mathbf{Pr}\left[\mathbf{X}_i = \max_{j \in S_e} \mathbf{X}_j\right]. \tag{1}$$

For the joint cost functions under discussion (equal to γ_e for every non-empty set), Definition 3.1 coincides with the definition given in Section 2. Since the value $\Delta_i(\pi)$ is γ_e for the first player of π and 0 otherwise, the equivalence follows.

Weighted Shapley shares are always increasing in a player's weight. If a set S_e contains at most two players, then the cost shares of Definition 3.1 are proportional to the players' weights. This is not generally true with three or more players.

Example 3.2. Suppose $\gamma_e = 1$ and $S_e = \{1, 2\}$ with $w_1 = 1$ and $w_2 = 2$. Since the edge has unit cost, the weighted Shapley share of player 1 is the probability that 1 is first in the random ordering described in Section 2. Hence it is equal to the probability that 2 is picked in the first sampling step, which gives us

$$\chi_{1,e}(\{1,2\}) = \frac{\frac{1}{w_2}}{\frac{1}{w_1}+\frac{1}{w_2}} = \frac{1}{3}.$$

Similarly, $\chi_{2,e}(\{1,2\}) = 2/3$, and the cost shares are proportional to the players' weights. Now suppose that player 3 with $w_3 = 1$ joins edge e. The weighted Shapley share of 1 is again the probability that 1 is first in the random ordering. This is now

$$\chi_{1,e}(\{1,2,3\}) = \frac{\frac{1}{w_2}}{\frac{1}{w_1}+\frac{1}{w_2}+\frac{1}{w_3}} \cdot \frac{\frac{1}{w_3}}{\frac{1}{w_1}+\frac{1}{w_3}} + \frac{\frac{1}{w_3}}{\frac{1}{w_1}+\frac{1}{w_2}+\frac{1}{w_3}} \cdot \frac{\frac{1}{w_2}}{\frac{1}{w_1}+\frac{1}{w_2}} = \frac{7}{30}.$$

Since $w_3 = w_1$, we also have $\chi_{3,e}(\{1,2,3\}) = 7/30$, and then $\chi_{2,e}(\{1,2,3\}) = 1-2\cdot 7/30 = 8/15$. These cost shares are not proportional to the players' weights.

We next show that every SV network cost-sharing game with the cost shares of Definition 3.1 admits a (weighted) potential function. Define the function Φ by

$$\Phi(P) = \sum_{e\in E} \Phi_e(P), \tag{2}$$

where the edge potential Φ_e is defined as

$$\Phi_e(P) = \gamma_e \cdot \mathbf{E}\left[\max_{j\in S_e} \mathbf{X}_j\right].$$

PROPOSITION 3.3. *For every pair P and $P' = (P_{-i}, P'_i)$ of path vectors that differ only in the ith component,*

$$\Phi(P') - \Phi(P) = w_i \cdot \left(C_i(P') - C_i(P)\right), \tag{3}$$

where C_i denotes the sum of the cost shares paid by player i.

PROOF. We prove that every edge contributes the same amount to the left- and right-hand sides of (3). If $e \in P_i \cap P'_i$ or $e \notin P_i \cup P'_i$, there is nothing to prove. By symmetry, we can assume that $e \in P'_i \setminus P_i$. We need to show that

$$\Phi_e(P') - \Phi_e(P) = w_i \cdot \chi_{i,e}(S_e \cup \{i\}), \tag{4}$$

where S_e is the set of players that use e in P.

The left-hand side of (4) is the difference between

$$\Phi_e(P') = \gamma_e \cdot \mathbf{E}\left[\max_{j\in S_e\cup\{i\}} \mathbf{X}_j\right] \text{ and } \Phi_e(P) = \gamma_e \cdot \mathbf{E}\left[\max_{j\in S_e} \mathbf{X}_j\right].$$

The maxima inside the expectations are different only when $\mathbf{X}_i$ is larger than the corresponding random variable of every player of S_e. Conditioning on this event and using the fact that the exponential distribution is memoryless, the conditional expected difference between the two maxima is $1/\lambda_i = w_i$. Hence $\Phi_e(P') - \Phi_e(P) = w_i \cdot \gamma_e \cdot \mathbf{Pr}\left[\mathbf{X}_i = \max_{j\in S_e\cup\{i\}} \mathbf{X}_j\right] = w_i \cdot \chi_{i,e}(S_e \cup \{i\})$, as claimed. □

As in Rosenthal [Rosenthal 1973a] and Monderer and Shapley [Monderer and Shapley 1996], the existence of a weighted potential function has immediate consequences. First, by (3), the outcome with minimum potential function value is a PNE. Moreover, every iteration of better-response dynamics — in which a player switches strategies to strictly decrease its cost — strictly decreases the potential function. Thus, better-response dynamics converges, necessarily to a PNE.

COROLLARY 3.4. *In every SV network cost-sharing game, better-response dynamics converges to a PNE.*

3.2. SV Atomic Selfish Routing.

In a *SV atomic selfish routing game*, each player $i = 1, 2, \ldots, k$ has a weight w_i and selects a path P_i from a node s_i to a node t_i in a given graph $G = (V, E)$. For every edge $e \in E$, the per-unit cost function $c_e(\cdot)$ is nonnegative and nondecreasing. Its users S_e have to pay a joint cost of

$$C_e(S_e) = x_e \cdot c_e(x_e), \tag{5}$$

where x_e is their total weight.

The joint cost function (5) is asymmetric, meaning that its value depends on the identities of the players in the set S_e and not just on $|S_e|$. This in is a contrast with weighted network cost-sharing games, where the asymmetry was exogenous to the (symmetric) joint cost function. For this reason, the standard (unweighted) Shapley value already gives meaningful weight-dependent cost shares in routing game with non-uniform player weights, and these are the cost shares proposed below. That is, we take the sampling parameter λ_i from Section 2 to be 1 for every player i (and not $1/w_i$). Section 3.3 outlines a natural generalization that accommodates both asymmetric cost functions and exogenous player asymmetry.

Definition 3.5. In an SV atomic selfish routing game, the Shapley share of player $i \in S_e$ on edge e is

$$\chi_{i,e}(S_e) = \mathbf{E}\left[C_e(S_e^i(\pi_e) \cup \{i\}) - C_e(S_e^i(\pi_e))\right],$$

where $S_e^i(\pi_e)$ denotes the players preceding i in π_e, a uniformly random ordering of S_e.

The cost shares in Definition 3.5 are generally proportional to players' weights if and only if the per-unit cost function c_e is affine.

Example 3.6. Suppose $c_e(x) = x$ and $S_e = \{1, 2\}$ with $w_1 = 1$ and $w_2 = 2$. Then the joint cost that the players have to share is $(w_1+w_2)^2 = 9$. The Shapley share of player 1 is

$$\chi_{1,e}(\{1,2\}) = \frac{1}{2} \cdot w_1^2 + \frac{1}{2} \cdot \left((w_1 + w_2)^2 - w_2^2\right) = 3.$$

Similarly we get $\chi_{2,e} = 6$ and see that the cost shares are proportional. Now suppose that $c_e(x) = x^2$ and S_e remains the same. The joint cost is $(w_1 + w_2)^3 = 27$ and

$$\chi_{1,e}(\{1,2\}) = \frac{1}{2} \cdot w_1^3 + \frac{1}{2} \cdot \left((w_1 + w_2)^3 - w_2^3\right) = 10, \quad \text{and}$$
$$\chi_{2,e}(\{1,2\}) = \frac{1}{2} \cdot w_2^3 + \frac{1}{2} \cdot \left((w_1 + w_2)^3 - w_1^3\right) = 17;$$

thus, the cost shares are not proportional.

Define a function Φ by

$$\Phi(P) = \sum_{e \in E} \Phi_e(P),$$

where the edge potential Φ_e is defined as

$$\Phi_e(P) = \sum_{i \in S_e} \chi_{i,e}(S_e^i(\pi) \cup \{i\}) \tag{6}$$

for some ordering π on S_e. For this definition to make sense, it must be the case that the right-hand side of (6) is independent of the ordering π. This is a special case of a result of Hart and Mas-Colell [Hart and Mas-Colell 1989] (see Section 3.3), for which we give a direct proof.

PROPOSITION 3.7. *For every joint cost function C with player set S, the value of*

$$\sum_{i \in S} \mathbf{E}_{\tau^i}\left[C(S^i(\pi,\tau^i) \cup \{i\}) - C(S^i(\pi,\tau^i))\right] \tag{7}$$

is the same for every ordering π of S, where τ^i is a permutation of $S^i(\pi) \cup \{i\}$ chosen uniformly at random and $S^i(\pi,\tau^i)$ denotes the players of S that precede i in both π and τ^i.

PROOF. For a fixed ordering π of the players, the quantity in (7) can be written as a sum of the form $\sum_{T \subseteq S} a_T c(T)$ for some set $\{a_T\}_{T \subseteq S}$ of coefficients. We now explicitly compute these coefficients and show that they do not depend on π.

Fix a subset $T \subseteq S$. With respect to the ordering π, let i denote the last player of T, say in position ℓ. There is a positive contribution to the coefficient a_T from the ℓth summand of (7), and a negative contribution from all subsequent summands. The positive contribution equals the probability that, among all random orderings of the players of $S^i(\pi) \cup \{i\}$, the players of T come first and player i is the last of these. This probability is

$$\frac{(|T|-1)! \cdot (\ell - |T|)!}{\ell!}.$$

Let i_j denote the jth player in the ordering π for some $j > \ell$. The negative contribution to the coefficient a_T by the jth summand of (7) equals the probability that, among all random orderings of the first j players under π, the players of T come first and are immediately followed by player i_j. This probability is

$$\frac{|T|! \cdot (j - |T| - 1)!}{j!}.$$

Summing over all players $j > \ell$ after T under π and rewriting, we obtain

$$a_T = (|T|-1)! \cdot \left[\frac{1}{\ell \cdot (\ell-1) \cdots (\ell - |T| + 1)} - \sum_{j=\ell+1}^{k} \frac{|T|}{j \cdot (j-1) \cdots (j - |T|)}\right], \tag{8}$$

where k is the number of players in S. Since

$$\frac{|T|}{j \cdot (j-1) \cdots (j-|T|)} = \left(\frac{1}{(j-1) \cdots (j-|T|)} - \frac{1}{j \cdots (j - |T| + 1)}\right)$$

for every $j > \ell$, the sum in (8) telescopes and hence

$$a_T = \frac{(|T|-1)!}{k \cdot (k-1) \cdots (k - |T| + 1)},$$

which is a function only of the sizes k and $|T|$ and is independent of the position ℓ of the final element of T in π. We conclude that the sum (7) is the same for every ordering π of the players in S. $\square$

The fact that the function Φ is a potential function now follows easily.

PROPOSITION 3.8. *For every pair P and $P' = (P_{-i}, P'_i)$ of path vectors that differ only in the ith component,*

$$\Phi(P') - \Phi(P) = C_i(P') - C_i(P). \tag{9}$$

PROOF. As in the proof of Proposition 3.3, we can focus on a single edge $e \in P'_i \setminus P_i$. By Proposition 3.7, we can compute the contribution of e to the left-hand side of (9) using an ordering π of the players in which i follows all of the players of S_e. Then, edge e contributes exactly $\chi_{i,e}(S_e \cup \{i\})$ to both sides of (9). □

COROLLARY 3.9. *In every SV atomic selfish routing game, better-response dynamics converges to a PNE.*

3.3. Arbitrary Congestion-Type Games

The (weighted) Shapley shares in Definitions 3.1 and 3.5 can be generalized to arbitrary congestion-type games. Consider a resource set E and a player set $S = \{1, 2, \ldots, k\}$, where each resource e has a joint cost functions $C_e : 2^S \to \mathbb{R}$ defined on the subsets of S, and each player i has a strategy set $\mathcal{P}_i \subseteq 2^E$ and a positive weight w_i. For a resource e, subset S_e of players, and a player $i \in S_e$, define the weighted Shapley share $\chi_{i,e}(S_e)$ of i for resource e when its users are S_e as its weighted Shapley value (Section 2) in the game with player set S_e and cost function C_e restricted to 2^{S_e}. The cost $C_i(P)$ to a player i in a strategy profile P is then defined as the sum of its cost shares:

$$C_i(P) = \sum_{e \in P_i} \chi_{i,e}(S_e),$$

where $S_e = \{j \in S : e \in P_j\}$ denotes the users of resource e in the profile P.

We claim that every game defined in this way admits a weighted potential function and hence better-response dynamics converges to a PNE. The argument follows that in Section 3.2. Define a function $\Phi = \sum_{e \in E} \Phi_e(P)$ in which the edge potential Φ_e is defined as

$$\Phi_e(P) = \sum_{i \in S_e} w_i \cdot \chi_{i,e}(S_e^i(\pi) \cup \{i\}) \tag{10}$$

for some ordering π on the players S_e using e in P. Hart and Mas-Colell [Hart and Mas-Colell 1989] proved that the right-hand side of (10) is independent of the ordering π, for every joint cost functions C_e and positive weight vector w. The proof that Φ is a weighted potential function is the same as in the proof of Proposition 3.8.

4. THE PRICE OF STABILITY IN SV NETWORK COST-SHARING GAMES.

This section provides tight bounds on the price of stability in SV network cost-sharing games — the ratio between the cost of the best PNE and the minimum-cost outcome. It is easy to see that, for every weight vector with k players, the worst PNE of such a game can cost k times as much as an optimal solution, and that this is tight [Chen et al. 2010].

Section 4.1 generalizes a construction of Anshelevich et al. [Anshelevich et al. 2008] to players with general weights. Section 4.2 is the primary contribution of this section, a matching upper bound for every positive weight vector.

4.1. POS Lower Bound.

Consider a graph $G = (V, E)$ and players $i = 1, 2, \ldots, k$ with distinct source vertices $s_1, \ldots, s_k$ and a common sink vertex t; see also Figure 2. As usual, we assume that $w_1 \leq w_2 \leq \cdots \leq w_k$. The graph has one additional vertex v. There is an edge

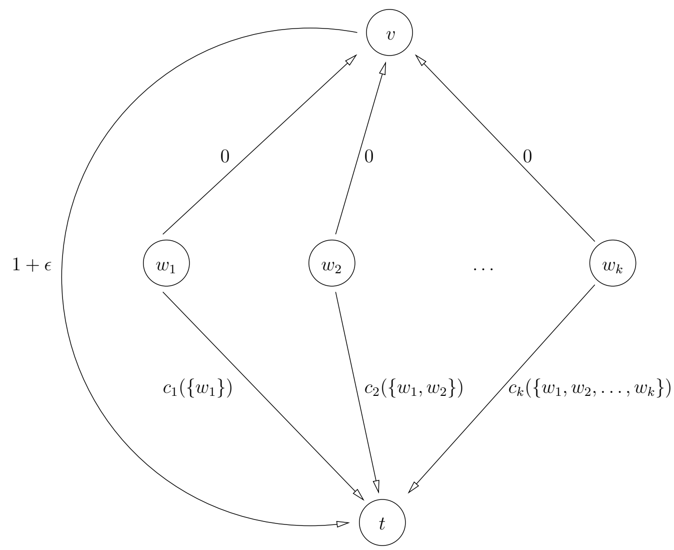

Fig. 2. Proof of Proposition 4.1. The worst-case price of stability is at least the expression in (11).

from v to t with cost $1+\epsilon$, where $\epsilon > 0$ is arbitrarily small. For each i, there is a zero-cost edge from s_i to v. For each i, the edge from s_i to t is set to the weighted Shapley share of the ith player for a unit-cost edge shared by players with weights $w_1, w_2, \dots, w_i$; we denote the quantity by $c_i(\{w_1, w_2, \dots, w_i\})$.

In the graph G, each player i can either use the path $s_i \to v \to t$, or use the direct edge from s_i to t. The optimal solution, in which every player i chooses the path $s_i \to v \to t$, has cost $1+\epsilon$. We claim that in the unique PNE of this SV network cost-sharing game, every player i chooses the direct s_i-t edge. To see this, consider the player k with the largest weight. The smallest cost it could have by taking the two-hop path is $(1+\epsilon) \cdot c_k(\{w_1, w_2, \dots, w_k\})$, which occurs when all players share the edge from v to t. This is larger than the cost of its one-hop path. Hence, in every PNE, player k uses its one-hop path and does not share the edge from v to t. The same reasoning applies inductively, showing that in every PNE, every player uses its one-hop path. This construction gives the following lower bound for every positive weight vector w.

PROPOSITION 4.1. *For every set of k players with positive nondecreasing weight vector w, the worst-case price of stability in SV network cost-sharing games with weight vector w is at least*

$$\sum_{i=1}^{k} c_i(\{w_1, w_2, \dots, w_i\}). \tag{11}$$

Setting $w = (1, 1, \dots, 1)$ recovers the well-known lower bound of $\mathcal{H}_k$ on the price of stability with unweighted players [Anshelevich et al. 2008].

4.2. POS Upper Bound.

The goal of this section is to prove that the lower bound in Proposition 4.1 is tight for every weight vector w.

THEOREM 4.2. *For every SV network cost-sharing game with player set $S = \{1, 2, \ldots, k\}$ and positive nondecreasing weight vector w, the price of stability is at most*

$$\sum_{i=1}^{k} c_i(\{w_1, w_2, \ldots, w_i\}). \tag{12}$$

The special case of unweighted players, where the bound (12) is the kth Harmonic number $\mathcal{H}_k$, has a short proof: the potential function Φ in (2) is always at least and never more than $\mathcal{H}_k$ times the cost of an outcome, so the potential function minimizer (a PNE) has cost at most $\mathcal{H}_k$ times that of an optimal outcome. With weighted players, the weighted potential function (2) need not approximate the cost of an outcome to any non-trivial factor, and a different argument is called for.

The high-level plan is as follows. We consider a minimum-cost outcome P^* and the outcome P that minimizes the weighted potential function Φ (2). To bound the cost of P in terms of P^*, we transform P^* into P one component at a time, in decreasing order of player weight. The change in outcome cost is the change in the deviating player's cost, which we can bound using the weighted potential function, plus the change in other players' cost. We argue that the worst case occurs when the deviating player abandons all edges it was using previously and switches only to edges that were previously unused. Bounding the cost of this worst case yields the theorem.

Before proceeding to the formal proof of Theorem 4.2, we prove a technical lemma. It states that the upper bound in (12) is nondecreasing in the player set. This is not obvious, as deleting a player removes one summand from (12) but also increases the value of some of the remaining summands.

LEMMA 4.3. *For every set $S = \{1, 2, \ldots, k\}$ of players with nondecreasing positive weight vector w, and every player $j \in S$,*

$$\sum_{i=1}^{k} c_i(\{w_1, w_2, \ldots, w_i\}) \geq \sum_{i=1}^{j-1} c_i(\{w_1, w_2, \ldots, w_i\}) + \sum_{i=j+1}^{k} c_i(\{w_1, w_2, \ldots, w_{j-1}, w_{j+1}, \ldots, w_i\}). \tag{13}$$

PROOF. (Proof of Lemma 4.3.) The first $j-1$ summands on both sides are the same. Only the left-hand side of (13) has a summand with $i = j$, namely $c_j(\{w_1, w_2, \ldots, w_j\})$. For $i > j$, the ith summand on the left-hand side ($c_i(\{w_1, w_2, \ldots, w_i\})$) is smaller than the corresponding summand on the right-hand side ($c_i(\{w_1, w_2, \ldots, w_{j-1}, w_{j+1}, \ldots, w_i\})$). To calculate the difference, we turn to the probabilistic representation of weighted Shapley shares in terms of exponentially distributed random variables $\mathbf{X}_1, \ldots, \mathbf{X}_k$ (Section 3.1).

In the right-hand side summand, the random variable $\mathbf{X}_i$ does not have to compete with $\mathbf{X}_j$ in order to be the largest. Hence, the event that $\mathbf{X}_i$ is smaller than $\mathbf{X}_j$ but larger than every other player's random variable contributes to $c_i(\{w_1, w_2, \ldots, w_{j-1}, w_{j+1}, \ldots, w_i\})$ but not to $c_i(\{w_1, w_2, \ldots, w_i\})$. We denote this probability by $p_j(i)$. Recalling the density and distribution functions of exponentially distributed random variables, we have

$$p_j(i) = \int_0^\infty \lambda_i \cdot e^{-\lambda_i \cdot x} \cdot e^{-\lambda_j \cdot x} \cdot \prod_{l=1}^{j-1} \left(1 - e^{-\lambda_l \cdot x}\right) \cdot \prod_{l=j+1}^{i-1} \left(1 - e^{-\lambda_l \cdot x}\right) dx.$$

Recalling that $w_1 \leq w_2 \leq \cdots \leq w_k$ and hence $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k$, the difference Δ between the left-hand and right-hand sides of (13) is

$$\begin{aligned}
\Delta &= c_j(\{w_1, w_2, \ldots, w_j\}) - \sum_{i=j+1}^{k} p_j(i) && (14)\\
&= \int_0^\infty \lambda_j \cdot e^{-\lambda_j \cdot x} \cdot \prod_{l=1}^{j-1}(1 - e^{-\lambda_l \cdot x}) - \sum_{i=j+1}^{k} \lambda_i \cdot e^{-(\lambda_i+\lambda_j)\cdot x} \cdot \prod_{l=1, l\neq j}^{i-1}(1 - e^{-\lambda_l \cdot x})dx && (15)\\
&\geq \int_0^\infty \lambda_j \cdot e^{-\lambda_j \cdot x} \cdot \prod_{l=1}^{j-1}(1 - e^{-\lambda_l \cdot x}) - \sum_{i=j+1}^{k} \lambda_j \cdot e^{-2\cdot\lambda_j \cdot x} \cdot \prod_{l=1, l\neq j}^{i-1}(1 - e^{-\lambda_l \cdot x})dx && (16)\\
&\geq \int_0^\infty \lambda_j \cdot e^{-\lambda_j \cdot x} \cdot \prod_{l=1}^{j-1}(1 - e^{-\lambda_l \cdot x}) \cdot \left[1 - \sum_{i=j+1}^{k} e^{-\lambda_j \cdot x} \cdot \prod_{h=j+1}^{i-1}(1 - e^{-\lambda_h \cdot x})\right] dx && (17)\\
&\geq \int_0^\infty \lambda_j \cdot e^{-\lambda_j \cdot x} \prod_{l=1}^{j-1}(1 - e^{-\lambda_l \cdot x}) \cdot \left[1 - \underbrace{e^{-\lambda_j \cdot x} \cdot \sum_{i=j+1}^{k}(1 - e^{-\lambda_j \cdot x})^{i-j-1}}_{\leq e^{-\lambda_j \cdot x} \cdot e^{\lambda_j \cdot x} = 1}\right] dx && (18)\\
&\geq 0. && (19)
\end{aligned}$$

Here, (14) and (15) follow by definition (of Δ, the weighted Shapley value $c_j(\cdot)$, and $p_j(i)$), (16) follows by the monotonicity of $\lambda \cdot e^{\lambda \cdot x}$ as a function of λ and by the fact that the values of λ are nonincreasing, (17) follows by taking out the common factor, (18) follows by evaluating the last product, and (19) as explained below the comment. This concludes the proof of the lemma. $\square$

We now prove Theorem 4.2

PROOF. (Proof of Theorem 4.2.) Let P^* and P denote a minimum-cost outcome and an outcome that minimizes the weighted potential function Φ in (2), respectively. For the analysis, we imagine each player i deviating from P_i^* to P_i in nonincreasing weight order, i.e., in the order $k, k-1, \ldots, 1$. Let T_e^i denote the players using edge e before player i switches strategies, and let $\Delta\Phi_i$ denote the change in Φ when i switches strategies. By Proposition 3.3, the change in player i's cost is exactly $(\Delta\Phi_i)/w_i$. To compute the change in other players' costs, recall that the sum of the weighted Shapley shares of an edge used by at least one player always equals the cost of that edge. Thus, for every edge $e \in P_i^* \setminus P_i$ with $|T_e^i| \geq 2$, player i's withdrawal from edge e increases the sum of the cost shares of the players of $T_e^i \setminus \{i\}$ by $\chi_{i,e}(T_e^i)$. Symmetrically, for every edge $e \in P_i \setminus P_i^*$, player i's arrival to edge e decreases the sum of cost shares of players in T_e^i (if any) by $\chi_{i,e}(T_e^i \cup \{i\})$. Overall, when player i switches from P_i^* to P_i, the outcome cost increases by at most

$$\frac{\Delta\Phi_i}{w_i} + \sum_{e \in P_i^* \setminus P_i \,:\, |T_e^i| \geq 2} \chi_{i,e}(T_e^i).$$

Summing over all players i, we obtain

$$C(P) - C(P^*) \leq \sum_{i=1}^{k} \frac{\Delta\Phi_i}{w_i} + \sum_{i=1}^{k} \sum_{e \in P_i^* \setminus P_i \,:\, |T_e^i| \geq 2} \chi_{i,e}(T_e^i). \quad (20)$$

To bound the first term of the right-hand side of (20), write

$$\sum_{i=1}^{k} \frac{\Delta\Phi_i}{w_i} = \underbrace{\frac{1}{w_k} \cdot \sum_{j=1}^{k} \Delta\Phi_j}_{\leq 0} + \sum_{i=1}^{k-1} \left(\underbrace{\frac{1}{w_i} - \frac{1}{w_{i+1}}}_{\geq 0} \right) \cdot \underbrace{\sum_{j=1}^{i} \Delta\Phi_j}_{\leq 0} \leq 0. \tag{21}$$

Since $w_1 \leq w_2 \leq \cdots \leq w_k$, every term $(\frac{1}{w_i} - \frac{1}{w_{i+1}})$ is nonnegative. Every term $\sum_{j=1}^{i} \Delta\Phi_j$ is the total potential function change of a sequence of moves that terminates in the outcome P that minimizes Φ, and hence is nonpositive. We conclude that the term $\sum_i \frac{\Delta\Phi_i}{w_i}$ in (20) is nonpositive.

We next upper bound the contribution of each edge e to the second term in (20). Let S_e^* denote the set of players that use e in P^*, and $S_e^i = S_e^* \cap \{1, 2, \ldots, i\}$. We claim that

$$\sum_{i\,:\,e \in P_i^* \setminus P_i, |T_e^i| \geq 2} \chi_{i,e}(T_e^i) \leq \sum_{i \in S_e^*\,:\,|S_e^i| \geq 2} \chi_{i,e}(S_e^i). \tag{22}$$

The right-hand side of (22) corresponds to the scenario in which every user of e in P^* abandons e when switching to its strategy in P.

Inequality (22) follows from three observations. First, for each $\ell = 2, 3, \ldots, |S_e^*|$, the right-hand side of (22) contains exactly one summand $\chi_{i,e}(S_e^i)$ in which $|S_e^i| = \ell$. The corresponding set S_e^i contains the ℓ lowest-indexed — and hence lowest-weight — players of S_e^*, of which i has maximum weight. Second, for each $\ell = 2, 3, \ldots, |S_e^*|$, the left-hand side of (22) contains at most one summand $\chi_{j,e}(T_e^j)$ in which $|T_e^j| = \ell$. The corresponding set T_e^j contains $h \geq 0$ players with index higher than j, who have already deviated to another path that contains e, and the $\ell - h$ lowest-weight players of S_e^*, of which j has maximum weight. Third, Definition 3.1 implies that the weighted Shapley share of a player is increasing in its own weight and decreasing in other players' weights. Thus, for every $\ell = 2, 3, \ldots, |S_e^*|$, the summand on the right-hand side of (22) with $|S_e^i| = \ell$ is at least the summand on the left-hand side with $|T_e^j| = \ell$ (if any).

We get

$$\begin{aligned} C(P) - C(P^*) &\leq \sum_{e \in E} \sum_{i \in S_e^*\,:\,|S_e^i| \geq 2} \chi_{i,e}(S_e^i) && (23) \\ &= \sum_{e \in E\,:\,S_e^* \neq \emptyset} \left(\sum_{i \in S_e^*} \chi_{i,e}(S_e^i) - \gamma_e \right) && (24) \\ &\leq \sum_{e \in E\,:\,S_e^* \neq \emptyset} \gamma_e \cdot \left(\sum_{i=1}^{k} c_i(w_1, w_2, \ldots, w_i) - 1 \right), && (25) \end{aligned}$$

where (23) follows from (20), (21), and (22), (24) is given by a simple transformation, and (25) follows by applying Lemma 4.3. Then, by the fact that $C(P^*) = \sum_{e\,:\,S_e^* \neq \emptyset} \gamma_e$,

$$C(P) \leq C(P^*) \cdot \sum_{i=1}^{k} c_i(w_1, w_2, \ldots, w_i),$$

which proves the theorem. □

As a special case, if players' weights are within a constant factor of each other, then $c_i(w_1, w_2, \ldots, w_i) = \Theta(1/i)$ for every i and the POS is $O(\log k)$. In contrast, when PNE exist under proportional cost shares, the POS in this case can be $\Theta(k)$ [Chen and Roughgarden 2009].

More generally, the POS bound in Proposition 4.1 and Theorem 4.2 approaches k as the players' weights become more dramatically spread out. For example, when $w_i = i$ for $i = 1, 2, \ldots, k$, calculations show that the POS is $O(\sqrt{k})$.

5. THE PRICE OF ANARCHY IN SV ATOMIC SELFISH ROUTING GAMES.

This section gives matching upper and lower bounds on the worst-case price of anarchy in SV atomic selfish routing games. Section 5.1 covers preliminaries. Section 5.2 proves a POA upper bound that is parameterized by the set of resource cost functions. Section 5.3 evaluates this upper bound for cost functions that are polynomials with nonnegative coefficients. Section 5.4 gives a construction showing that, for every set of cost functions satisfying some mild technical conditions, this POA upper bound is tight in the worst case.

5.1. Preliminaries.

The worst-case POA in SV atomic selfish routing games depends on the set of allowable cost functions. For example, with cost functions that are polynomials with degree at most d and nonnegative coefficients, we prove that the worst-case POA is exponential in d, but independent of the network size and the number of players. This dependence motivates parameterizing our POA bounds by the class $\mathcal{C}$ of allowable resource cost functions. We do not expect the worst-case POA to admit a closed-form expression for every set $\mathcal{C}$, and instead seek a relatively simple characterization of this value. Throughout this section, we make the following assumptions.

(1) Every cost function $c \in \mathcal{C}$ is nonnegative, nondecreasing, and convex.
(2) The set $\mathcal{C}$ is closed under scaling and dilation, meaning that if $c(x) \in \mathcal{C}$ and $a, b > 0$, then $a \cdot c(b \cdot x) \in \mathcal{C}$.

5.2. POA Upper Bound.

Our upper bound approach is an instantiation of the "smoothness framework" articulated in [Roughgarden 2009]. Let $\xi^c(x, x^*)$ denote the worst-case Shapley share of a player with weight x^* on a resource with cost function c, when the total weight of the other players on the resource is x. This worst-case is taken over all possible sets of players with total weight x. We call a pair (λ, μ) of real numbers *feasible* for a cost function c if $\mu < 1$ and if

$$\xi^c(x^*, x) \leq \lambda \cdot x^* \cdot c(x^*) + \mu \cdot x \cdot c(x) \tag{26}$$

for every $x, x^* \geq 0$. We use $\mathcal{A}(\mathcal{C})$ to denote the set of pairs (λ, μ) that are feasible for every cost function $c \in \mathcal{C}$. Define

$$\zeta(\mathcal{C}) = \inf_{(\lambda,\mu)\in\mathcal{A}(\mathcal{C})} \frac{\lambda}{1-\mu},$$

or as $+\infty$ if $\mathcal{A}(\mathcal{C}) = \emptyset$.

Under the assumptions described in Section 5.1, $\zeta(\mathcal{C})$ is an upper bound on the POA of every SV atomic selfish routing game with cost functions in $\mathcal{C}$.

THEOREM 5.1. *Let $\mathcal{C}$ be a set of nonnegative, nondecreasing, and convex cost functions. Then the POA of every SV atomic selfish routing game with cost functions in $\mathcal{C}$ is at most $\zeta(\mathcal{C})$.*

PROOF. Let P and P^* denote a PNE and an arbitrary outcome of such a routing game with players $S = \{1, 2, \ldots, k\}$. Since P is a PNE, we have

$$\begin{aligned} C(P) &= \sum_{i=1}^{k} \sum_{e \in P_i} \chi_{i,e}(S_e) \\ &\leq \sum_{i=1}^{k} \sum_{e \in P_i^*} \chi_{i,e}(S_e \cup \{i\}), \end{aligned} \tag{27}$$

where S_e denotes the players using edge e in P. By the definition of $\xi^c(x^*, x)$,

$$\chi_{i,e}(S_e \cup \{i\}) \leq \xi^{c_e}(w_i, x_e), \tag{28}$$

where w_i is the weight of player i and x_e the total weight of the players in S_e. Then we get

$$C(P) \leq \sum_{i=1}^{k} \sum_{e \in P_i^*} \xi^{c_e}(w_i, x_e) \leq \sum_{e \in P_i^*} \xi^{c_e}(x_e^*, x_e),$$

where x_e^* denotes the total weight of players using edge e in P^*. The second inequality follows from the fact that c_e is convex, with the following reasoning: Given that c_e is convex and increasing, it follows that, for every x, the function $(w+x)\cdot c_e(w+x) - x \cdot c_e(x)$ is convex in w. Now consider any subset of e's users $T \subseteq S_e$ such that $\sum_{i \in T} w_i = x$ and suppose there is a set of players T^* such that $\sum_{i \in T^*} w_i = x_e^*$. If we place each $i \in T^*$, on resource e with the other users being T, we get a marginal increase $\delta_i = (w_i + x) \cdot c_e(w_i + x) - x \cdot c_e(x)$. Since, as we argued, this is a convex expression in w_i, the sum $\sum_{i \in T^*} \delta_i$ is maximized when there is only one player in T^*. The Shapley value is a sum over all possible orderings, hence over several possible sets $T \subseteq S_e$, which implies that the sum of the Shapley values is again maximized when there is a single player in T^*, which concludes the proof of our claim.

Now choose $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$. Since (λ, μ) satisfies (26) for all $c \in \mathcal{C}$ and $x, x^* \geq 0$, we have

$$\begin{aligned} C(P) &\leq \sum_{e \in E} \xi^{c_e}(x_e^*, x_e) \\ &\leq \sum_{e \in E} [\lambda \cdot x_e^* \cdot c_e(x_e^*) + \mu \cdot x_e \cdot c_e(x_e)] \\ &= \lambda \cdot C(P^*) + \mu \cdot C(P); \end{aligned}$$

rearranging terms completes the proof. □

Remark 5.2. The POA upper bound in Theorem 5.1 is a "smoothness proof" in the sense of Roughgarden [Roughgarden 2009]. Informally, this means that the hypothesis that P is a PNE is used only in the inequality (27), with hypothetical deviations P_i^* that are independent of the choice of P. This fact is interesting because POA bounds that are proved with smoothness arguments extend automatically to numerous other equilibrium concepts. Specifically, the POA upper bound of $\zeta(\mathcal{C})$ applies more generally to mixed-strategy Nash equilibria, correlated equilibria, and outcome sequences generated by no-regret learners [Roughgarden 2009]. Approximate Nash equilibria and polynomial-length best-response sequences also approximately obey the $\zeta(\mathcal{C})$ bound [Roughgarden 2009]. Finally, the POA bound of $\zeta(\mathcal{C})$ extends to all Bayes-Nash equilibria of incomplete information SV selfish routing games, where players' weights and source-sink pairs are drawn from an arbitrary product prior distribution [Roughgarden 2012; Syrgkanis 2012].

5.3. Example: Polynomial Cost Functions.

This section explicitly evaluates the POA upper bound in Theorem 5.1 for the special case in which $\mathcal{C}$ is the set of polynomials with nonnegative coefficients and maximum degree d.

Elementary calculus shows that, for every positive integer d, the function

$$g_d(x) = 3 \cdot x^{d+1} - 1 - (x+1)^{d+1} \tag{29}$$

has a unique positive root, which we denote by χ_d. This section establishes the following theorem.

THEOREM 5.3. *If $\mathcal{C}$ is the set of polynomials with nonnegative coefficients and maximum degree d, then the POA of a SV atomic selfish routing game with cost functions in $\mathcal{C}$ is at most $\chi_d^{d+1} = (\Theta(d))^{d+1}$.*

Remark 5.7 shows that the bound of χ_d^{d+1} is tight in the worst case, for every positive integer d. For comparison, the worst-case POA with proportional (rather than SV) cost-sharing, in such games that happen to possess PNE, is the slightly smaller quantity $\Theta((d/\ln d)^{d+1})$.

Before presenting the proof of Theorem 5.3, we examine the asymptotic behavior of χ_d.

PROPOSITION 5.4. *As $d \to \infty$, $\chi_d = \Theta(d)$.*

PROOF. Note that

$$g_d(d) = 3 \cdot d^{d+1} - 1 - (d+1)^{d+1} = 3 \cdot d^{d+1} - 1 - d^{d+1} \cdot (1+1/d) \cdot (1+1/d)^d.$$

Similarly,

$$g_d(d/2) = 3 \cdot (d/2)^{d+1} - 1 - (d/2)^{d+1} \cdot (1+2/d) \cdot \left((1+2/d)^{d/2}\right)^2.$$

Since $\lim_{x\to\infty}(1+1/x)^x = e$,

$$\lim_{d\to\infty} g_d(d) > 0 \text{ and } \lim_{d\to\infty} g_d(d/2) < 0.$$

Since $g_d(x)$ is decreasing up to point $1/(3^{1/d}-1)$ and increasing after that, it follows that its root χ_d is indeed between $d/2$ and d, for all sufficiently large d. More careful computations of this type show that χ_d tends to infinity as roughly $d/\ln 3 \approx 0.9 \cdot d$. $\square$

We now prove Theorem 5.3.

PROOF. (Proof of Theorem 5.3.) We exhibit values (λ, μ) that are feasible for every cost function $c \in \mathcal{C}$ — recall (26) — and that satisfy $\lambda/(1-\mu) \leq \chi_d^{d+1}$. The theorem then follows from Theorem 5.1.

Define

$$\lambda_j = \frac{(\chi_j+1)^j+1}{2} \text{ and } \mu_j = \frac{\left(\chi_j^{-1}+1\right)^j-1}{2} \text{ for } j = 1, 2, \ldots, d,$$

$\lambda = \max_j \lambda_j$, and $\mu = \max_j \mu_j$. It is evident that $\lambda = \lambda_d$.

We begin by showing that (λ, μ) is feasible. To see that $\mu < 1$, recall that, by definition,

$$3 \cdot \chi_j^{j+1} = 1 + (\chi_j+1)^{j+1}. \tag{30}$$

If $\mu_j \geq 1$, then $(\chi_j^{-1}+1)^j \geq 3$, which implies that $(1+\chi_j)^j \geq 3\cdot\chi_j^j$ and hence $(1+\chi_j)^{j+1} \geq 3\cdot\chi_j^j + 3\cdot\chi_j^{j+1}$. Combining this with (30) yields the contradiction $3\cdot\chi_j^j \leq -1$.

To see that (λ,μ) satisfies (26) for every cost function $c_e \in \mathcal{C}$, fix a specific resource e, a set of users S_e, and a player i, outside S_e. By Definition 3.5,

$$\chi_{i,e}(S_e \cup \{i\}) = \mathbf{E}[(\mathbf{X}_{i,e}+w_i)\cdot c_e(\mathbf{X}_{i,e}+w_i) - \mathbf{X}_{i,e}\cdot c_e(\mathbf{X}_{i,e})], \tag{31}$$

where w_i is the weight of player i and $\mathbf{X}_{i,e}$ is the total weight of the players preceding i in a uniformly random ordering of the players in $S_e \cup \{i\}$.

Let x_e denote the total weight of the players in S_e. Pairing up subsets of S_e with their complements, the right-hand side of (31) is a convex combination of terms of the form $\frac{1}{2}[(z+w_i)\cdot c_e(z+w_i)-z\cdot c_e(z)]+\frac{1}{2}[((x_e-z)+w_i)\cdot c_e((x_e-z)+w_i)-(x_e-z)\cdot c_e(x_e-z)]$. Since $c_e(x)$ is a polynomial with nonnegative coefficients, the term $(x+w_i)\cdot c_e(x+w_i) - x\cdot c_e(x)$ is convex and nondecreasing in x. It, then, follows that each of these terms is maximized when $z = x_e$. Thus,

$$\begin{aligned} C(P) &\leq \sum_{i=1}^{k}\sum_{e\in P_i^*}\left[\frac{1}{2}((x_e+w_i)\cdot c_e(x_e+w_i) - x_e\cdot c_e(x_e)) + \frac{1}{2}\cdot w_i\cdot c_e(w_i)\right] \\ &\leq \sum_{e\in E}\left[\frac{1}{2}((x_e+x_e^*)\cdot c_e(x_e+x_e^*) - x_e\cdot c_e(x_e)) + \frac{1}{2}\cdot x_e^*\cdot c_e(x_e^*)\right], \end{aligned}$$

where x_e^* denotes the total weight of players using edge e in P^*, with the second inequality following from the fact that the function $(x+w)\cdot c(x+w) - x\cdot c(x)$ is superadditive in w for every fixed x.

By linearity, the condition (26) reduces to proving that

$$\frac{(x+x^*)^{j+1}}{2} - \frac{x^{j+1}}{2} + \frac{(x^*)^{j+1}}{2} \leq \lambda\cdot(x^*)^{j+1} + \mu\cdot x^{j+1}. \tag{32}$$

for every $j = 1,2,\ldots,d$ and $x, x^* \geq 0$. Every λ_n, μ_n pair clearly satisfies inequality (32) when $x^* = 0$. Assume that $x^* > 0$ and set $r = x/x^*$. Rewrite inequality (32) as

$$(2\cdot\mu_j+1)\cdot r^{j+1} - (1+r)^{j+1} + (2\cdot\lambda_j - 1) \geq 0, \text{ for all } r \geq 0. \tag{33}$$

Considering the left-hand side of (33) as a function of r and taking the derivative, we can see that the minimizer is $r = \left((2\mu+1)^{1/j}-1\right)^{-1} = \chi_j$. With these values of r, λ_j, μ_j, the left-hand side of (33) equals 0, which verifies inequality (33) (and (32)). Inequality (32) clearly remains valid for $\lambda \geq \lambda_j$ and $\mu \geq \mu_j$, and so (λ,μ) form a feasible pair.

To prove that $\lambda/(1-\mu) \leq \chi_d^{d+1}$, recall that $\lambda = \lambda_d$ and write $\mu = \mu_\ell$ for some $\ell \in \{1,2,\ldots,d\}$. Then,

$$\begin{aligned} \frac{\lambda}{1-\mu} &= \frac{(\chi_d+1)^d+1}{3-(\chi_\ell^{-1}+1)^\ell} \\ &= \frac{1}{3}\cdot\frac{3\cdot\chi_\ell^{\ell+1}\cdot\left((\chi_d+1)^d+1\right)}{3\cdot\chi_\ell^{\ell+1} - \chi_\ell\cdot(\chi_\ell+1)^\ell} \\ &= \frac{1}{3}\cdot\frac{(\chi_\ell+1)^{\ell+1}+1}{(\chi_\ell+1)^\ell+1}\cdot\left((\chi_d+1)^d+1\right), \end{aligned}$$

where the last step follows from (30). The last expression is clearly increasing in ℓ. Hence, setting $\ell = d$ and using (30) once again, we derive $\lambda/(1-\mu) \leq \chi_d^{d+1}$, as required. $\square$

5.4. POA Lower Bound.

The upper bounds presented in Section 5.2 are tight in the worst case. The construction that proves this is simplest to present in the context of general congestion games where players have strategy sets that are arbitrary subsets of the edges and not necessarily paths. It is easy to convert the construction into an atomic selfish routing network.

THEOREM 5.5. *For every class $\mathcal{C}$ that is closed under scaling and dilation, the POA of a SV atomic congestion game with cost functions in $\mathcal{C}$ can be arbitrarily close to $\zeta(\mathcal{C})$.*

Our construction resembles one used previously to prove POA lower bounds for weighted congestion games with proportional cost shares [Bhawalkar et al. 2010], but some of the technical details differ. Our proof of Theorem 5.5 requires the following technical lemma.

LEMMA 5.6. *Let $\mathcal{C}$ be a class of cost functions with $\zeta(\mathcal{C}) > 1$. For every positive $\epsilon < \zeta(\mathcal{C}) - 1$, at least one of the following conditions holds.*

(*1*) *There exist $c \in \mathcal{C}, x \geq 0, x^* > 0$ such that*

$$\xi^c(x^*, x) \geq x \cdot c(x) \;\; \textit{and} \;\; \frac{x \cdot c(x)}{x^* \cdot c(x^*)} \geq \zeta(\mathcal{C}) - \epsilon.$$

(*2*) *There exist $c_1, c_2 \in \mathcal{C}$, $x_1, x_2 \geq 0$, $x_1^*, x_2^* > 0$, and λ, μ such that*

$$\begin{aligned}
\xi^{c_1}(x_1^*, x_1) &= \lambda \cdot x_1^* \cdot c_1(x_1^*) + \mu \cdot x_1 \cdot c_1(x_1); \\
\xi^{c_2}(x_2^*, x_2) &= \lambda \cdot x_2^* \cdot c_2(x_2^*) + \mu \cdot x_2 \cdot c_2(x_2); \\
\xi^{c_1}(x_1^*, x_1) &\leq x_1 \cdot c_1(x_1); \\
\xi^{c_2}(x_2^*, x_2) &\geq x_2 \cdot c_2(x_2); \\
\textit{and} \;\; \frac{\lambda}{1-\mu} &> \zeta(\mathcal{C}) - \epsilon.
\end{aligned}$$

The intuition behind the lemma and the subsequent construction is as follows. The optimal (λ, μ) lies on the two-dimensional space and within a polytope. It will either be very close to a single constraint or to the intersection of two constraints. Then we can use that single constraint or a convex combination of the two intersecting constraints to construct a lower bound with every player (other than the leaf players) facing a situation equivalent to what the constraints describe.

As an example, for the case with a single constraint, each player faces the equivalent of the following situation: The player has weight x^* and can deviate to a resource with cost function c where there are already players with total weight x (the actual distribution of the weight x among the various players is as in the worst case for the player with weight x^*). To obtain a tight bound we will have to make every player indifferent between the equilibrium strategy and the potential deviation. After a very large number of layers of such players, the construction is terminated by a final layer of players such that each one is indifferent between the equilibrium strategy and a currently empty resource of equal cost. This means all the players in the final layer have the same cost in the optimal solution and in the equilibrium, so their contribution to the total cost must be negligible for the POA bound to be tight. For this to be true, the cost of players as layers progress must not increase, which translates to the condition $\xi^c(x^*, x) \geq x \cdot c(x)$, since $\xi^c(x^*, x)$ is the equilibrium cost of the x^* player (given that, as mentioned, the players are indifferent between their equilibrium strategies and their potential deviations) and the overall cost of the player's competitors in the lower layer is $x \cdot c(x)$.

For the case when two constraints intersect near the optimal (λ, μ) pair, the corresponding c_1, x_1, x_1^* and c_2, x_2, x_2^* are combined in a way that is equivalent to having a single c, x, x^* which satisfies the "correct lower bound conditions" we described in the previous paragraph.

The proof of Lemma 5.6 geometrically describes why we can always find the necessary constraints (lines in the two-dimensional space) and the proof of Theorem 5.5 describes the details of the lower bound construction that uses these constraints.

PROOF. (Proof of Lema 5.6) For a cost function $c \in \mathcal{C}$, $x \geq 0$, and $x^* > 0$, let $\mathcal{H}_{c,x,x^*}$ denote the half-plane

$$\xi^c(x^*, x) \leq \lambda \cdot x^* \cdot c(x^*) + \mu \cdot x \cdot c(x)$$

and $\partial\mathcal{H}_{c,x,x^*}$ the boundary of this half-plane. Recall from (26) that these are the half-planes that define the set $\mathcal{A}(\mathcal{C})$ of feasible pairs (λ, μ) for the set $\mathcal{C}$ of cost functions. Also, define

$$\beta_{c,x,x^*} = \frac{x \cdot c(x)}{\xi^c(x^*, x)}$$

and

$$\zeta_{c,x,x^*} = \frac{x \cdot c(x)}{x^* \cdot c(x^*)}.$$

Fix a positive $\epsilon < \zeta(\mathcal{C}) - 1$ and let $\zeta' = \zeta(\mathcal{C}) - \epsilon/2$. If $\zeta(\mathcal{C})$ is not finite, set $\zeta' = 1/\epsilon$. We write $\mathcal{L}_{\zeta'}$ for the line $\lambda + \zeta' \cdot \mu = \zeta'$ in the λ, μ plane.

If we think of a boundary line $\partial\mathcal{H}_{c,x,x^*}$ as specifying μ as a function of λ, then this line has slope $-1/\zeta_{c,x,x^*}$ and μ-intercept $1/\beta_{c,x,x^*}$. The half-space $\mathcal{H}_{c,x,x^*}$ consists of everything "northeast" of its boundary.

Consider the half-planes with $\beta_{c,x,x^*} \leq 1$. In the lucky event that there is such a half-plane with $\zeta_{c,x,x^*} \geq \zeta'$, we are done: this choice of c, x, x^* satisfies the first set of conditions of the lemma. For the rest of the proof, we assume that $\zeta_{c,x,x^*} < \zeta'$ for every half-plane with $\beta_{c,x,x^*} \leq 1$.

We consider two cases. To define them, pick an arbitrary cost function c_1 with $c_1(1) > 0$ — since $\mathcal{C}$ is closed under dilation, such a function exists — and a sufficiently large value of x_1 so that $\zeta_{c_1,x_1,1} > \zeta'$. Our standing assumption implies that $\beta_{c_1,x_1,1} > 1$. Define $(\hat{\lambda}, \hat{\mu})$ as the unique point of intersection of $\partial H_{c_1,x_1,1}$ and $\mathcal{L}_{\zeta'}$. Since the former line has a larger slope ($-1/\zeta_{c_1,x_1,1}$ vs. $-1/\zeta'$) and a smaller μ-intercept ($1/\beta_{c_1,x_1,1}$ vs. 1) than the latter, $\hat{\lambda} > 0$ and hence $\hat{\mu} < 1$.

For case (i), we assume that there exists a half-plane $\mathcal{H}_{c_2,x_2,x_2^*}$ with $\beta_{c_2,x_2,x_2^*} < 1$ whose boundary intersects the line $\mathcal{L}_{\zeta'}$ at a point (λ_2, μ_2) with $\mu_2 < \hat{\mu}$. Equivalently, the line $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersects $\mathcal{L}_{\zeta'}$ to the right of where $\partial\mathcal{H}_{c_1,x_1,1}$ intersects $\mathcal{L}_{\zeta'}$. Since the μ-intercepts of $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ and $\partial\mathcal{H}_{c_1,x_1,1}$ (namely, $1/\beta_{c_2,x_2,x_2^*} > 1$ and $1/\beta_{c_1,x_1,1} < 1$) are on either side of that of $L_{\zeta'}$ (namely, 1) and $\hat{\lambda} > 0$, this implies that the intersection (λ, μ) of $\partial\mathcal{H}_{c_1,x_1,1}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ lies on the "northeast side" of $L_{\zeta'}$. It follows that $\lambda + \zeta' \cdot \mu \geq \zeta'$. Thus, $c_1, c_2, x_1, x_2, 1, x_2^*, \lambda, \mu$ satisfy the conditions in the second set of conditions of the lemma. This argument is illustrated in Figure 3.

Finally, assume that all half-planes $\mathcal{H}_{c,x,x^*}$ with $\beta_{c,x,x^*} < 1$ have boundaries that intersect the line $\mathcal{L}_{\zeta'}$ at points (λ, μ) with $\mu \geq \hat{\mu}$. This is handled by case (ii), which is illustrated in Figure 4. Let μ^* denote the infimum of all μ-coordinates of such intersections. Under our standing assumption, every such boundary $\partial\mathcal{H}_{c,x,x^*}$ has a smaller slope ($-1/\zeta_{c,x,x^*}$ vs. $-1/\zeta'$) and a larger μ-intercept ($1/\beta_{c_1,x_1,1}$ vs. 1) than $\mathcal{L}_{\zeta'}$, and hence intersects $\mathcal{L}_{\zeta'}$ at a point (λ, μ) with $1 > \mu \geq \hat{\mu}$. Thus, $1 > \mu^* \geq \hat{\mu}$.

We now find appropriate (c_1, x_1, x_1^*) and (c_2, x_2, x_2^*) with $\beta_{c_1,x_1,x_1^*} \geq 1$ and $\beta_{c_2,x_2,x_2^*} < 1$, such that the corresponding half-plane boundaries intersect $\mathcal{L}_{\zeta'}$ at points (λ_1, μ_1)

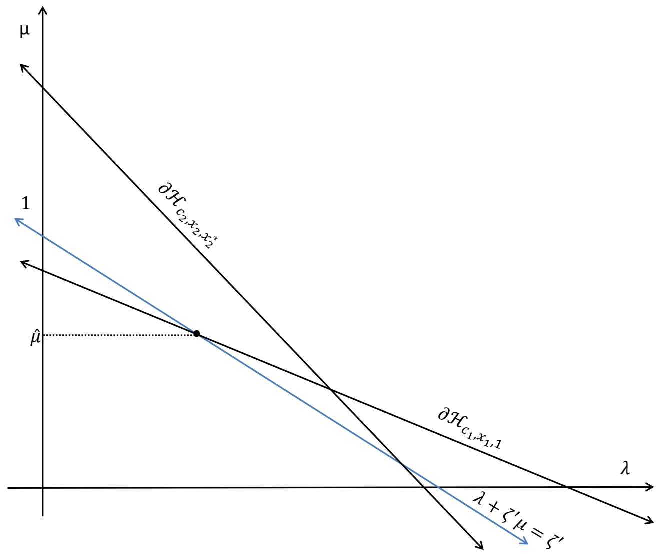

Fig. 3. Case (i) from the proof of Lemma 5.6: there exists (c_2, x_2, x_2^*) with $\beta_{c_2,x_2,x_2^*} < 1$ that intersects the line $\mathcal{L}_{\zeta'}$ at a point λ, μ with $\mu < \hat{\mu}$.

and (λ_2, μ_2) with μ_1, μ_2 very close to μ^*. Let $\delta = \frac{\epsilon\cdot(1-\mu^*)}{4\cdot\zeta'-\epsilon} > 0$. Consider the point $(\zeta'\cdot(1-\mu^*+\delta), \mu^*-\delta)$ of $\mathcal{L}_{\zeta'}$. This point is feasible for all constraints corresponding to (c, x, x^*) with $\beta_{c,x,x^*} < 1$. Since $\zeta' < \zeta(\mathcal{C})$, this point cannot belong to the feasible set $\mathcal{A}(\mathcal{C})$ and hence there exists (c_1, x_1, x_1^*) with $\beta_{c_1,x_1,x_1^*} \geq 1$ such that the point $(\zeta'\cdot(1-\mu^*+\delta), \mu^*-\delta)$ violates the corresponding constraint. Note that the point $(0, 1)$ of $\mathcal{L}_{\zeta'}$ lies in $\mathcal{H}_{c_1,x_1,x_1^*}$. This implies that $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ intersects $\mathcal{L}_{\zeta'}$ at a point (λ_1, μ_1) with $\mu_1 \geq \mu^*-\delta$. Moreover, $\lambda_1 + \zeta'\cdot\mu_1 = \zeta'$.

If $\mu_1 > \mu^*$, then we can find (c_2, x_2, x_2^*) with $\beta_{c_2,x_2,x_2^*} < 1$ that intersects $\mathcal{L}_{\zeta'}$ at (λ_2, μ_2) with $\mu^* \leq \mu_2 \leq \mu_1$. Then, similarly to the previous case, $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersect at a point (λ, μ) such that $\lambda/(1-\mu) \geq \zeta'$, completing the proof.

We can now assume that $\mu^* - \delta \leq \mu_1 \leq \mu^*$. By the definition of μ^*, there exist (c_2, x_2, x_2^*) such that $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersects $\mathcal{L}_{\zeta'}$ at (λ_2, μ_2), with $\mu^* \leq \mu_2 \leq \mu^* + \delta$. Note that $\mu_2 \geq \mu_1$ and $\lambda_2 + \zeta'\cdot\mu_2 = \zeta'$.

Let (λ, μ) be the point where $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersect. Both these boundaries have negative slopes, which means (λ, μ) lies in the triangle formed by the points (λ_1, μ_1), (λ_2, μ_2), and (λ_2, μ_1). Then $\lambda/(1-\mu) \geq \lambda_2/(1-\mu_1)$. Since $\lambda_1 - \lambda_2 = \zeta'\cdot(\mu_2-\mu_1) \leq 2\cdot\zeta'\cdot\delta$, we have

$$\begin{aligned}\frac{\lambda_2}{1-\mu_1} &= \frac{\lambda_1}{1-\mu_1} - \frac{\lambda_1-\lambda_2}{1-\mu_1}\\ &\geq \zeta' - \frac{2\cdot\zeta'\cdot\delta}{1-\mu^*+\delta}\\ &\geq \zeta' - \frac{\epsilon}{2}.\end{aligned}$$

This proves that the second set of conditions in the statement of the lemma hold. □

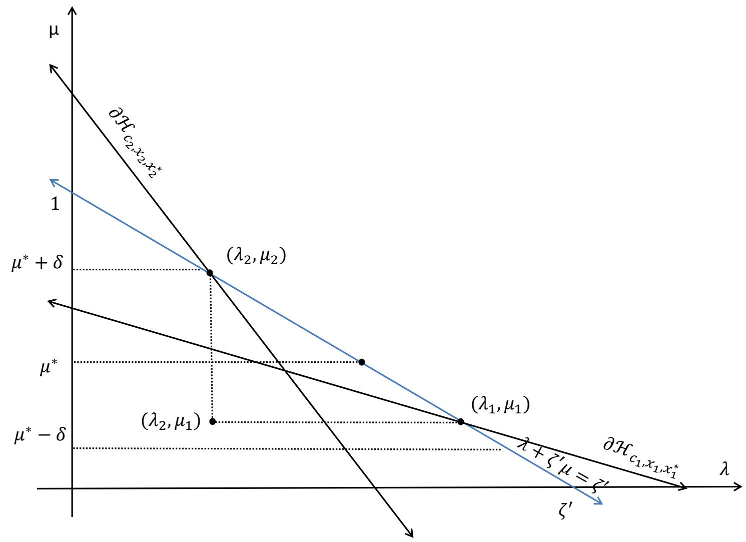

Fig. 4. Case (ii) from the proof of Lemma 5.6: all $\mathcal{H}_{c,x,x^*}$ with $\beta_{c,x,x^*} < 1$ have boundaries that intersect $\mathcal{L}_{\zeta'}$ at (λ,μ) with $\mu \geq \hat{\mu}$.

Before proceeding with the proof of Theorem 5.5, we take note of some consequences of Lemma 5.6. Suppose the second case of the lemma applies and offers two triples (c_1, x_1, x_1^*) and (c_2, x_2, x_2^*) such that the corresponding half-plane boundaries intersect at (λ, μ) with $\lambda/(1-\mu) > \zeta(\mathcal{C}) - \epsilon$. Scaling and dilating a cost function does not affect the corresponding constraint (26). Thus, for every $w > 0$, we can find cost functions $\hat{c}_1$ and $\hat{c}_2$ such that

$$\begin{aligned}\xi^{\hat{c}_1}(w, w\cdot z_1) &= \lambda\cdot w\cdot \hat{c}_1(w) + \mu\cdot w\cdot z_1\cdot \hat{c}_1(w\cdot z_1);\\ \xi^{\hat{c}_2}(w, w\cdot z_2) &= \lambda\cdot w\cdot \hat{c}_2(w) + \mu\cdot w\cdot z_2\cdot \hat{c}_2(w\cdot z_2),\end{aligned} \tag{34}$$

where $z_1 = x_1/x_1^*$ and $z_2 = x_2/x_2^*$.

Moreover, since $\xi^{\hat{c}_1}(w, w\cdot z_1) \leq z_1\cdot \hat{c}_1(z_1\cdot w)$ and $\xi^{\hat{c}_2}(w, w\cdot z_2) \geq z_2\cdot \hat{c}_2(z_2\cdot w)$, there is a constant $\eta \in [0,1]$ such that

$$\begin{aligned}\eta\cdot z_1\cdot \hat{c}_1(w\cdot z_1) + (1-\eta)\cdot z_2\cdot \hat{c}_2(w\cdot z_2) \;=\; &\eta\cdot \xi^{\hat{c}_1}(w, w\cdot z_1)\\ &+(1-\eta)\cdot \xi^{\hat{c}_2}(w, w\cdot z_2).\end{aligned} \tag{35}$$

We now give our lower bound construction.

PROOF. (Proof of Theorem 5.5.) Our proof has two cases, corresponding to the two sets of conditions of Lemma 5.6. First consider a set $\mathcal{C}$ and $\epsilon > 0$ so that the second set of conditions of the lemma applies and the corresponding triples are (c_1, x_1, x_1^*) and (c_2, x_2, x_2^*). We write S_{c_1,x_1,x_1^*} for the set of players with total weight x_1 that yield the Shapley share of $\xi^{c_1}(x_1^*, x_1)$ for the player with weight x_1^*. Similarly, we write S_{c_2,x_2,x_2^*} for the corresponding set of players of the second triple.

We now describe our game, the equilibrium strategies, and the potential deviations of the players in S. The construction is organized into m layers of resources. The 1-st layer has a single resource and a single player using that resource in our proposed equilibrium, P. This player has weight 1 and the cost function of the resource is $c \in \mathcal{C}$

such that $c(1) = 1$ (such a function clearly exists, since $\mathcal{C}$ is closed under dilation). We now recursively present the following layers, up to the $(m-1)$-th layer.

Consider any player that is using a resource on the i-th layer, with $i < m-1$, in the proposed equilibrium P. Suppose this player has weight w_i and cost c_i in P. There is only one other alternative strategy for the player, which is to play a pair of resources, l, r, from the $(i+1)$-th layer. There is a unique l, r pair of resources in the $(i+1)$-th layer for every player in the i-th layer. We now pick the cost functions for l and r and the set of players using l and r in P. We start with the sets of users in P. The users of l are $|S_{c_1,x_1,x_1^*}|$ and their weights are precisely those in S_{c_1,x_1,x_1^*} scaled by w_i/x_1^*. Similarly, the users of r are $|S_{c_2,x_2,x_2^*}|$ and their weights are precisely those in S_{c_2,x_2,x_2^*} scaled by w_i/x_2^*. Hence the total weight on l in P is $w_i \cdot x_1/x_1^* = w_i \cdot z_1$ and the total weight on r in P is $w_i \cdot x_2/x_2^* = w_i \cdot z_2$. The users of l and r are disjoint and use no other resources in P. Among all pairs of cost functions that satisfy (34) for $x^* = w_i$, pick a pair that also satisfies

$$c_j(w_i \cdot z_j) \cdot z_j = c_i \tag{36}$$

for $j \in \{1, 2\}$. Since $\mathcal{C}$ is closed under scaling and dilation, such a pair exists. Let η_i be the corresponding value for η in (35) and define

$$c_l(x) = \eta_i \cdot c_1(x) \text{ and } c_r(x) = (1 - \eta_i) \cdot c_2(x). \tag{37}$$

We finally present the m-th layer. Consider any player who uses a resource on the $(m-1)$-th layer in P. Let w_i be the weight and c_i the cost of the player in P. The only alternative strategy of the player is to play a resource on the m-th layer with cost function $c \in \mathcal{C}$ such that$c(w_i) = c_i$ (since $\mathcal{C}$ is closed under scaling and dilation, such a function exists). There are no players that use resources of the m-th layer in P.

We claim that the POA of the above game is $\lambda/(1-\mu)$, where λ, μ are the parameters in the second guarantee of Lemma 5.6. We first prove that P is a PNE. It is clear that a player on the $(m-1)$-th layer has no incentive to deviate, since the corresponding deviation to the m-th layer would yield the exact same cost. Consider any other player i who uses resource e and has cost c_i. Let $\{l, r\}$ be the alternative strategy on the next layer. Then, using (36) first and (35), (37) subsequently, we get

$$\begin{aligned}
\chi_{i,e}(\{i\}) &= c_i = w_i \cdot \eta_i \cdot z_1 \cdot c_1(w_i \cdot z_1) + w_i \cdot (1 - \eta_i) \cdot z_2 \cdot c_2(w_i \cdot z_2) \\
&= w_i \cdot \left[\frac{1}{2} \cdot c_r(w_i) + \frac{1}{2} \cdot (1 + z_2) \cdot c_r(w_i \cdot (1 + z_2)) - \frac{1}{2} \cdot z_2 \cdot c_r(w_i \cdot z_2)\right] \\
&\quad + w_i \cdot \left[\frac{1}{2} \cdot c_l(w_i) + \frac{1}{2} \cdot (1 + z_1) \cdot c_l(w_i \cdot (1 + z_1)) - \frac{1}{2} \cdot z_1 \cdot c_l(w_i \cdot z_1)\right] \\
&= \chi_{i,r}(S_r \cup \{i\}) + \chi_{i,l}(S_l \cup \{i\}).
\end{aligned}$$

This completes the proof that P is a PNE. Also, combining the second line of the above equality with (34), we get

$$\begin{aligned}
\chi_{i,e}(\{i\}) &= \lambda \cdot (\chi_{i,l}(\{i\}) + \chi_{i,r}(\{i\})) + \mu \cdot \chi_{i,e}(\{i\}) \text{ which gives} \\
\chi_{i,e}(\{i\}) &\geq (\zeta(\mathcal{C}) - \epsilon) \cdot (\chi_{i,l}(\{i\}) + \chi_{i,r}(\{i\})).
\end{aligned} \tag{38}$$

In the outcome P, the contribution of a player i, who is above layer $m-1$, to the total cost $C(P)$ is equal to the combined contributions of the players using the corresponding l and r resources. This follows from (37) and (36):

$$\begin{aligned}
z_1 \cdot w_i \cdot c_l(z_1 \cdot w_i) &+ z_2 \cdot w_i \cdot c_r(z_2 \cdot w_i) \\
&= z_1 \cdot w_i \cdot \eta_i \cdot c_1(z_1 \cdot w_i) + z_2 \cdot w_i \cdot (1 - \eta_i) \cdot c_r(z_2 \cdot w_i) \\
&= w_i \cdot c_e(w_i).
\end{aligned}$$

It follows that, in P, the combined contribution of each layer is the same. Since the contribution of the 1-st layer is equal to 1, we get that $C(P) = m$. Now, let P^* be the outcome that has all players play the other strategy than the one in P. Call S_m the set of players in the m-th layer in P^* and let $e(i)$ be the (only) resource that i uses in P.

$$\begin{aligned} C(P^*) &= \sum_{i \in S} w_i \cdot \sum_{e \in P_i^*} \chi_{i,e}(\{i\}) \\ &= \sum_{i \in S_m} w_i \cdot \sum_{e \in P_i^*} \chi_{i,e}(\{i\}) + \sum_{i \in S \setminus S_m} w_i \cdot \sum_{e \in P_i^*} \chi_{i,e}(\{i\}) \\ &= \sum_{i \in S_m} w_i \cdot \chi_{i,e(i)}(\{i\}) + \sum_{i \in S \setminus S_m} w_i \cdot \frac{\chi_{i,e(i)}(\{i\})}{\zeta(\mathcal{C}) - \epsilon} = 1 + \frac{m-1}{\zeta(\mathcal{C}) - \epsilon}, \end{aligned}$$

where the last line follows from (38). The fact that $\lim_{m\to\infty} C(P)/C(P^*) = \zeta(\mathcal{C}) - \epsilon$ concludes the first case of the proof.

Finally, consider a set $\mathcal{C}$ and $\epsilon > 0$ so that the first guarantee of Lemma 5.6 applies. Thus, there is a triple (c, x, x^*) with

$$\xi^c(x^*, x) \geq x \cdot c(x)$$

and

$$\frac{x \cdot c(x)}{x^* \cdot c(x^*)} \geq \zeta(\mathcal{C}) - \epsilon.$$

A similar construction yields the lower bound in this case. Let $z = x/x^*$. Then for $w > 0$ we have

$$\xi^c(w, w \cdot z) \geq z \cdot w \cdot c(z \cdot w)$$

and

$$\frac{z \cdot c(z \cdot w)}{c(w)} \geq \zeta(\mathcal{C}) - \epsilon.$$

Repeating the construction of the previous case with $(c_1, x_1, x_1^*) = (c, x, x^*)$, and $\eta_i = 1$ for every player i before the final layer, yields the same lower bound and the same analysis holds (the triple c_2, x_2, x_2^* is left unspecified since it is practically nonexistent given that $1 - \eta_i$ will always be 0). □

Remark 5.7. Here we illustrate this construction in the special case of cost functions that are polynomials with nonnegative coefficients and degree at most d. Note that for $c(x) = x^d$, and $x = \chi_d, x^* = 1$, the conditions for the first case of Lemma 5.6 hold (using the fact (30) that $3 \cdot \chi_d^{d+1} - 1 = (\chi_d + 1)^{d+1}$).

We have shown in the proof of Theorem 5.3 that the worst case for these cost functions is when there is a single player comprising the whole x. We can therefore infer that every layer of the construction will have only one player and that, as we move along the layers, the weights are scaled by a factor $x/x^* = \chi_d$. The weight of the ith player is χ_d^{i-1} and the cost function of the jth resource is $\chi_d^{(1-j)\cdot(d+1)} \cdot x^d$ — except for the last one, which has the same cost function as its neighbor. The POA in this example is χ_d^{d+1}, matching the upper bound in Theorem 5.3.

REFERENCES

ALAND, S., DUMRAUF, D., GAIRING, M., MONIEN, B., AND SCHOPPMANN, F. 2011. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing 40,* 5, 1211–1233.

ANSHELEVICH, E., DASGUPTA, A., KLEINBERG, J., TARDOS, É., WEXLER, T., AND ROUGHGARDEN, T. 2008. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing 38,* 4, 1602–1623.

AWERBUCH, B., AZAR, Y., AND EPSTEIN, L. 2005. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*. 57–66.

BHAWALKAR, K., GAIRING, M., AND ROUGHGARDEN, T. 2010. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*. Vol. 2. 17–28.

CHEN, H. AND ROUGHGARDEN, T. 2009. Network design with weighted players. *Theory of Computing Systems 45,* 2, 302–324.

CHEN, H., ROUGHGARDEN, T., AND VALIANT, G. 2010. Designing network protocols for good equilibria. *SIAM Journal on Computing 39,* 5, 1799–1832.

FOTAKIS, D., KONTOGIANNIS, S. C., AND SPIRAKIS, P. G. 2005. Selfish unsplittable flows. *Theoretical Computer Science 348,* 2-3, 226–239.

GAIRING, M. AND SCHOPPMANN, F. 2007. Total latency in singleton congestion games. In *Proceedings of the 7th International Workshop on Internet and Network Economies (WINE)*. 381–387.

GOEMANS, M. X., MIRROKNI, V., AND VETTA, A. 2005. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*. 142–151.

GOPALAKRISHNAN, R., MARDEN, J. R., AND WIERMAN, A. 2013. Potential games are necessary to ensure pure Nash equilibria in cost sharing games. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*. 563–564.

HARKS, T. AND KLIMM, M. 2012. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research 37,* 3, 419–436.

HARKS, T., KLIMM, M., AND MÖHRING, R. H. 2011. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems 49,* 1, 46–70.

HART, S. AND MAS-COLELL, A. 1989. Potential, value, and consistency. *Econometrica 57,* 3, 589–614.

KALAI, E. AND SAMET, D. 1987. On weighted Shapley values. *International Journal of Game Theory 16,* 3, 205–222.

KOUTSOUPIAS, E. AND PAPADIMITRIOU, C. H. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*. 404–413.

MARDEN, J. R. AND WIERMAN, A. 2013. Distributed welfare games. *Operations Research 61,* 1, 155–168.

MILCHTAICH, I. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior 13,* 1, 111–124.

MONDERER, D. AND SHAPLEY, L. S. 1996. Potential games. *Games and Economic Behavior 14,* 1, 124–143.

MOSK-AOYAMA, D. AND ROUGHGARDEN, T. 2009. Worst-case efficiency analysis of queueing disciplines. In *Proceedings of the 36th International Colloqium on Automata, Languages and Programming (ICALP)*. 546–557.

MOULIN, H. 2008. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory 36,* 3, 379–405.

OSBORNE, M. J. AND RUBINSTEIN, A. 1994. *A Course in Game Theory*. MIT Press.

ROSENTHAL, R. W. 1973a. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory 2,* 1, 65–67.

ROSENTHAL, R. W. 1973b. The network equilibrium problem in integers. *Networks 3,* 1, 53–59.

ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *41st ACM Symposium on Theory of Computing (STOC)*. 513–522.

ROUGHGARDEN, T. 2012. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 862–879.

ROUGHGARDEN, T. AND TARDOS, É. 2002. How bad is selfish routing? *Journal of the ACM 49,* 2, 236–259.

SHAPLEY, L. S. 1953. Additive and non-additive set functions. Ph.D. thesis, Department of Mathematics, Princeton University.

SHENKER, S. J. 1995. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking 3,* 6, 819–831.

SYRGKANIS, V. 2012. Bayesian games and the smoothness framework. arXiv.cs.GT:1203.5155v1.

Received ; revised ; accepted