Network Cost-Sharing without Anonymity*

Tim Roughgarden[†] Okke Schrijvers[‡]

July 18, 2014

Abstract

We consider network cost-sharing games with non-anonymous cost functions, where the cost of each edge is a submodular function of its users, and this cost is shared using the Shapley value. Non-anonymous cost functions model asymmetries between the players, which can arise from different bandwidth requirements, durations of use, services needed, and so on.

These games can possess multiple Nash equilibria of wildly varying quality. The goal of this paper is to identify well-motivated equilibrium refinements that admit good worst-case approximation bounds. Our primary results are tight bounds on the cost of strong Nash equilibria and potential function minimizers in network cost-sharing games with non-anonymous cost functions, parameterized by the set $\mathcal{C}$ of allowable submodular cost functions.

These two worst-case bounds coincide for every set $\mathcal{C}$, and equal the *summability* parameter introduced in [Roughgarden and Sundararajan, 2009] to characterize efficiency loss in a family of cost-sharing mechanisms. Thus, a single parameter simultaneously governs the worst-case inefficiency of network cost-sharing games (in two incomparable senses) and cost-sharing mechanisms. This parameter is always at most the kth Harmonic number $\mathcal{H}_k \approx \ln k$, where k is the number of players, and is constant for many function classes of interest.

1 Introduction

1.1 Network Cost-Sharing Games with Non-Anonymous Cost Functions

We consider network cost-sharing games with non-anonymous cost functions. Such a game takes place in a directed graph $G = (V, E)$ and has k players. Player i has a source vertex $s_i \in V$ and a sink vertex $t_i \in V$, and its strategy set is the s_i-t_i paths of the graph.[1] Outcomes of the game correspond to path vectors $(P_1, \ldots, P_k)$, with the semantics that the subnetwork $(V, \cup_{i=1}^{k} P_i)$ gets formed.

Each edge e has a cost function C_e, specifying the total cost incurred on edge e as a function of its users — the players S_e that pick a path that includes e. The function $C_e(S_e)$ models the

*An extended abstract of this paper appears in the Proceedings of the the 7th International Symposium on Algorithmic Game Theory, October 2014.

[†]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF grants CCF-1016885 and CCF-1215965, and an ONR PECASE Award. Email: `tim@cs.stanford.edu`.

[‡]Department of Computer Science, Stanford University, 482 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF grant CCF-1016885. Email: `okke@cs.stanford.edu`.

[1]The main results of this paper continue to hold, with the same proofs, when the strategy set of a player i is an arbitrary subset of 2^E.

infrastructure or service cost of supporting the users S_e between e's endpoints. We always assume that $C_e(\emptyset) = 0$ and that C_e is monotone, meaning $S_e \subseteq T_e$ implies $C_e(S_e) \leq C_e(T_e)$. For most of the paper, we assume that C_e is submodular, meaning it exhibits diminishing costs in the following sense:

$$C_e(T_e \cup \{i\}) - C_e(T_e) \leq C_e(S_e \cup \{i\}) - C_e(S_e)$$

for all i and $S_e \subseteq T_e$. Almost all previous work on network cost-sharing games, beginning with [Anshelevich et al., 2008a], considers only *anonymous* cost functions, where the cost of an edge depends solely on the number of users. For anonymous cost functions, submodularity is equivalent to non-increasing marginal costs.

Example 1.1 (Weighted Players) For a simple example of a non-anonymous cost function, suppose each player i has a positive weight w_i, with higher weights indicating greater demands for resources (more bandwidth, longer duration, etc.). Suppose that the joint cost function $C_e(S_e)$ depends on the set S_e of users only through the sum of their weights $\sum_{i \in S_e} w_i$. If $C_e(S_e) = f(\sum_{i \in S_e} w_i)$ for a nondecreasing concave function f, then C_e is monotone and submodular.

Example 1.2 (Coverage Functions) For a more general class of non-anonymous cost functions, consider a ground set X of services, where supporting a service $j \in X$ imposes a weight of w_j on the service provider. Each player i requires a set $A_i \subseteq X$ of services. The cost C_e of supporting all of the services required by a set S_e of users is $C_e(S_e) = f(\sum_{j \in \cup_{i \in S_e} A_i} w_j)$. Provided f is a monotone concave function, C_e is a monotone submodular function. Example 1.1 corresponds to the special case in which the players require disjoint sets of services. If all of the A_i's coincide, we recover the special case of (anonymous) constant cost functions.

To complete the description of the game, we need to define players' costs. To do this, we assign a *cost share* $\chi_e(i, S_e)$ to each user $i \in S_e$ of each edge e. The cost $c_i(\mathcal{P})$ of a player i in a strategy profile $\mathcal{P} = (P_1, \ldots, P_k)$ is then

$$c_i(\mathcal{P}) = \sum_{e \in P_i} \chi_e(i, S_e),$$

where $S_e = \{j : e \in P_j\}$ is the set of users of e.

With anonymous cost functions, the natural cost shares proposed in [Anshelevich et al., 2008a] are the equal cost shares: $\chi_e(i, S_e) = C_e(S_e)/|S_e|$. With non-anonymous cost functions, however, such cost shares are not as well motivated. We extend the idea of equal cost-sharing to non-anonymous cost functions by taking $\chi_e(i, S_e)$ to be i's *Shapley value* in the cooperative game induced by C_e and S_e. In more detail, for a permutation σ of the players of S_e, let $\Delta_\sigma(i)$ denote the increase $C(S_{\sigma(1..i-1)} \cup \{i\}) - C(S_{\sigma(1..i-1)})$ in cost due to i's arrival, where $S_{\sigma(1..i-1)}$ is the set of players that precede i in σ. Then, $\chi_e(i, S_e)$ is defined as the expected value of $\Delta_\sigma(i)$, where the expectation is over the (uniform at random) choice of σ. It is easy to verify that: (i) these cost shares coincide with equal cost-sharing when C_e is anonymous; (ii) the joint cost is shared fully across the players, with $\sum_{i \in S_e} \chi_e(i, S_e) = C_e(S_e)$; and (iii) a submodular cost function C_e leads to positive externalities, in the sense that $\chi_e(i, T_e) \leq \chi_e(i, S_e)$ whenever $i \in S_e \subseteq T_e$. Properties (ii) and (iii) are called *budget-balance* and *cross-monotonicity.* These cost shares also ensure that every network cost-sharing game has a potential function and therefore admits at least one pure Nash equilibrium (see Section 2).

Example 1.3 (Weighted Players Revisited) Consider two players with weights 1 and 3 and the cost function $C_e(S_e) = (\sum_{i \in S_e} w_i)^{1/2}$. The joint cost of the two players is 2. The players' cost shares are their Shapley values, namely $\frac{1}{2}(3 - \sqrt{3}) \approx .635$ and $\frac{1}{2}(1 + \sqrt{3}) \approx 1.365$, respectively.[2]

1.2 Measures of Equilibrium Inefficiency

The primary goal of this paper is to characterize the inefficiency of equilibria in network cost-sharing games with non-anonymous cost functions, in as many senses as possible. We next review the relevant measures of equilibrium inefficiency.

We define the cost $C(\mathcal{P})$ of a strategy profile $\mathcal{P}$ of a network cost-sharing game as the sum of players' costs:

$$C(\mathcal{P}) = \sum_{i=1}^{k} c_i(\mathcal{P}) = \sum_{e \in E} C_e(S_e), \tag{1}$$

and take (1) as our objective function. Recall that a pure Nash equilibrium (PNE) is a strategy profile from which no player can decrease its cost via a unilateral deviation. To what extent do PNE minimize the cost (1)?

1.2.1 A Non-Starter: The Price of Anarchy

It is well known that network cost-sharing games can have multiple PNE of wildly varying quality. The canonical example in the basic model with constant cost functions [Anshelevich et al., 2008a,b] posits k players, each choosing between an edge e_1 with fixed cost $1 + \epsilon$ and an edge e_2 with fixed cost k. Since all players using the second edge is a PNE, the ratio between the worst PNE and an optimal solution (i.e., the "price of anarchy") can be as large as k. A simple argument gives a matching upper bound on the worst-case price of anarchy.[3]

1.2.2 Equilibrium Refinements

The bad equilibrium identified above does not imply that network cost-sharing games are uninteresting — just that, to reason meaningfully about the quality of their equilibria, a more fine-grained approach is required. Recall that an *equilibrium refinement* defines a subset of equilibria. We are interested in equilibrium refinements with the following two properties:

1. All equilibria in the refined set have cost close to optimal.

[2] Non-uniform player weights have been proposed previously with anonymous cost functions [Anshelevich et al., 2008a, Chen and Roughgarden, 2009, Chen et al., 2010, Kollias and Roughgarden, 2011], necessarily in a different way. An anonymous cost function is by definition insensitive to the players' weights, and hence the Shapley cost shares are equal, independent of the weights. One approach to incorporating weights is to redefine the cost shares to be proportional to players' weights [Anshelevich et al., 2008a], at the cost of losing pure Nash equilibria and good efficiency guarantees for them [Chen and Roughgarden, 2009]. Another way to enforce weight-dependent cost shares is to switch to a weighted variant of the Shapley value [Kollias and Roughgarden, 2011]. In the present work, with (non-anonymous) cost functions that depend directly on players' weights, the standard Shapley value already leads to weight-dependent cost shares. As we show, these cost shares guarantee the existence of good pure Nash equilibria. Our model also captures more general player asymmetries that cannot be modeled simply through heterogeneous player weights.

[3] More generally, the price of anarchy depends on the allowable edge cost functions, but it remains disappointingly large (polynomial in k) for essentially all cost functions of interest.

2. There is a plausible explanation why equilibria in the refinement are more “important” or “likely” than those outside the set.

Previous work on network-cost sharing games with anonymous submodular cost functions can be interpreted as proposing two refinements with these two properties: potential function minimizers and strong Nash equilibria.

1.2.3 Potential Function Minimizers and the Price of Stability

Anshelevich et al. [2008a] proposed circumventing bad PNE by studying the price of stability, defined as the ratio between the minimum-cost PNE and that of an optimal outcome. They prove that the worst-case price of stability in network cost-sharing games with anonymous submodular cost functions is exactly the kth Harmonic number $\mathcal{H}_k = \sum_{i=1}^{k} \frac{1}{i} = \ln k + \Theta(1)$. The worst case is achieved by an instance with constant cost functions.

We can interpret the upper bound in [Anshelevich et al., 2008a] as a worst-case bound for an equilibrium refinement by examining its proof. The first step in [Anshelevich et al., 2008a] constructs a potential function [Monderer and Shapley, 1996] for every network cost-sharing game with anonymous cost functions — a function Φ such that the change in Φ under a unilateral deviation by player i equals the change in i's cost. The PNE correspond to the local minimizers (under unilateral deviations) of Φ. The second step of the proof shows that every global minimizer of Φ has cost at most $\mathcal{H}_k$ times that of an optimal outcome.[4] The global minimizers of Φ therefore form an equilibrium refinement of the PNE that satisfies the first property in Section 1.2.2.[5] As for the second property, global potential function optimizers have been previously proposed as a plausible equilibrium refinement in the game theory and economics literature, together with supporting theoretical [Asadpour and Saberi, 2009, Blume, 1993, Slade, 1994] and experimental evidence [Chen and Chen, 2011].

1.2.4 Strong Nash Equilibria

Epstein et al. [2009] studied an incomparable equilibrium refinement in network cost-sharing games with anonymous submodular cost functions. Recall that a strong Nash equilibrium (SNE) [Aumann, 1959] is a strategy profile such that no coalition of players can deviate in a coordinated way to strictly decrease all of their costs. Robustness to coalitional deviations provides good motivation for favoring strong Nash equilibria over non-strong PNE.

Global potential function minimizers need not be SNE — indeed, SNE are not guaranteed to exist in general [Epstein et al., 2009] — and SNE need not minimize the potential function, except in very simple networks [Holzman and Law-Yone, 1997]. Nevertheless, Epstein et al. [2009] proved that the worst-case ratio between an SNE and an optimal outcome is also precisely $\mathcal{H}_k$.

1.3 Contributions and Paper Organization

Our results are summarized in Table 1. Our primary contribution is a characterization of the worst-case inefficiency of both potential function minimizers and strong Nash equilibria in network

[4] The lowest-cost PNE need not be a global potential function minimizer; see [Tardos and Wexler, 2007] and Proposition 3.8 in the current work.

[5] The proof in [Anshelevich et al., 2008a] also shows that the PNE reachable from a minimum-cost outcome via better-response dynamics form such an equilibrium refinement. The results in the present work also hold for this alternative refinement.

Table 1: Summary of main results. The columns describe the existence of a potential function and bounds on the different solution concepts, and the rows give the conditions under which the results hold. "SM" stands for submodular. "Shapley" indicates Shapley cost shares. The parameter α denotes the summability (2) of the cost-sharing method for the cost function.

	Potential Φ	**Φ minimizer $= \alpha$**	**Strong PoA $= \alpha$**	**Seq PoA $= \alpha$**
Anonymous costs	✓	✓	✓	✓
SM & Shapley	✓	✓	✓	✗
Non-SM, Shapley	✓	$\leq \alpha$, not tight	✗	✗
Non-SM, cross-mono	✗	N/A	✓	✗

cost-sharing games, as a function of the class $\mathcal{C}$ of allowable submodular cost functions. These results correspond to the second and third columns of Table 1. Despite their instance-by-instance incomparability, we prove that the worst-case approximation ratios of these two equilibrium refinements are identical for every set $\mathcal{C}$ of allowable cost functions, and equal the worst-case *summability* of a function in $\mathcal{C}$.[6]

In more detail, we recall the summability parameter introduced in [Roughgarden and Sundararajan, 2009] to characterize efficiency loss in cost-sharing mechanisms [Moulin, 1999]. Let $C(\cdot)$ denote a cost function defined on a ground set U and χ a *cost-sharing method* – a function from player sets and players to cost shares. We call χ *α-summable for C*, if for every player subset $S \subseteq U$ and every ordering σ of the players of S,

$$\sum_{\ell=1}^{|S|} \chi(i_{\sigma(\ell)}, S_{\sigma(1..\ell)}) \leq \alpha \cdot C(S), \tag{2}$$

where $i_{\sigma(\ell)}$ denotes the ℓth player in σ, and $S_{\sigma(1..\ell)}$ denotes the set of the first ℓ players in the ordering σ. In words, we begin with the empty set and add players of S one-by-one according to σ. Letting X_ℓ denote the cost share of the ℓth player (according to χ) when the player is first added, the cost-sharing method χ is α-summable for C if the sum $\sum_\ell X_\ell$ only overestimates the cost of $C(S)$ by an α factor (for a worst-case choice of the subset S and the ordering σ of the players). Let $\alpha(C)$ denote the smallest α such that C is α-summable and $\alpha(\mathcal{C}) = \sup_{C \in \mathcal{C}} \alpha(C)$.

We prove in Section 3 that the worst-case ratio between the cost of a potential function minimizer or a strong Nash equilibrium and an optimal outcome of a network cost-sharing game with (monotone and submodular) cost functions in $\mathcal{C}$ is precisely $\alpha(\mathcal{C})$. This bound is always at most $\mathcal{H}_k$, where $k = |U|$ (Proposition 2.2), and is constant for many cost function classes of interest. For example, consider weighted players (Example 1.1) and a polynomial cost function $C_e(S_e) = (\sum_{i \in S_e} w_i)^d$. For $d \in (0, 1]$, this function is at most $\frac{1}{d}$-summable, independent of the number of players and their weights (Example 2.4). This yields a constant-factor approximation guarantee for potential function minimizers and strong Nash equilibria in games with such cost functions. Intuitively, the approximation guarantee is constant except for "approximately constant" cost functions.

In addition to our main results above, in Section 4 we extend the equivalence of summability and worst-case approximation bounds for strong Nash equilibria to network cost-sharing games

[6] Our upper bound for strong Nash equilibria meets the "coalitional smoothness" criterion of [Bachrach et al., 2013], and therefore extends to additional solution concepts.

with non-submodular cost functions and arbitrary cross-monotonic cost-sharing methods (fourth row of Table 1). At this level of generality the games typically have no potential function, so the first refinement is not well defined.

Our results in Sections 3 and 4 effectively isolate the key features of the standard network cost-sharing model that drive inefficiency bounds. While there exist properties that require the full symmetry of anonymous cost functions (see Section 1.4), tight worst-case bounds for potential function minimizers only rely on cross-monotonicity of the underlying cost-sharing method and (obviously) the existence of a potential function, and tight bounds for strong Nash equilibria do not even require a potential function.

Section 5 interprets known results for anonymous cost functions in terms of the summability parameter, and shows that these fail to carry over to non-anonymous functions. Section 6 clarifies the relationship between network cost-sharing games and Moulin mechanisms [Moulin, 1999] by showing a sense in which the latter are a special case of the former. Section 7 identifies conditions on the cost functions that extend the existence guarantees in [Epstein et al., 2009] for strong Nash equilibria in restricted network topologies beyond anonymous cost functions. Section 8 concludes.

1.4 Further Related Work

This paper contributes to the literature on network cost-sharing games that was initiated in [Anshelevich et al., 2008a,b]. Subsequent works on these and related models include [Fiat et al., 2006] [Li, 2009] [Bilò et al., 2010] [Christodoulou et al., 2010] [Bilò and Bove, 2011] [Kawase and Makino, 2013] [Lee and Ligett, 2013] [Chekuri et al., 2007] [Charikar et al., 2008] [Syrgkanis, 2010] [Chen and Roughgarden, 2009] [Chen et al., 2010] [Kollias and Roughgarden, 2011] [Paes Leme et al., 2012]. All of these papers study only anonymous cost functions. Most consider only the special case of constant cost functions (with $C_e(S_e) = \gamma_e$ if $S_e \neq \emptyset$ and 0 otherwise), with the goal of obtaining better efficiency guarantees in undirected and other special classes of networks. There are two previous papers that treat network cost-sharing games with non-anonymous cost functions. Gopalakrishnan et al. [2014] characterize cost-sharing rules that guarantee the existence of PNE and do not consider equilibrium inefficiency. Von Falkenhausen and Harks [2013] consider machine scheduling games that may have positive and negative externalities and design cost sharing methods that yield small PoA bounds. Their cost-sharing methods, unlike ours, are "non-uniform," and are defined in an instance-specific way that references the optimal outcome of the given instance.

Not all efficiency guarantees for network cost-sharing games with anonymous submodular cost functions can be extended to the non-anonymous case. For example, one result in [Paes Leme et al., 2012] can be interpreted as an approximation guarantee for a third equilibrium refinement for the case of network cost-sharing games with anonymous submodular cost functions and (asymmetric) singleton player strategies. They consider a sequential version of the game and show that (assuming no ties) all player orderings yield (permutations of) the same subgame perfect equilibrium, which in turn induces a pure Nash equilibrium of the single-shot game with near-optimal cost. With non-anonymous submodular cost functions and singleton strategies, there need not be such an "order-independent" subgame perfect equilibrium. Similarly, the algorithm in [Syrgkanis, 2010] for efficiently computing a near-optimal PNE in such games does not extend to non-anonymous submodular cost functions. See Section 5 for more details.

Many other models of network formation have been studied. We mention only the network creation game of [Fabrikant et al., 2003], which has been particularly influential in the theoretical computer science literature; see [Tardos and Wexler, 2007] for a survey. For a comprehensive treatment of game-theoretic models of network formation, see [Jackson, 2008].

The idea of using the Shapley value to define cost-sharing methods with good properties is not new; see [Moulin, 1999] and the references therein. In network cost-sharing games, the weighted Shapley value [Shapley, 1953, Kalai and Samet, 1987] was first introduced implicitly in [Chen et al., 2010] for games with constant cost functions to characterize cost-sharing rules that guarantee the existence of pure Nash equilibria. This characterization was recently generalized by [Gopalakrishnan et al., 2014] to all cost functions. Shapley value-based cost shares have also been used in congestion games, which can be thought of as the "negative externality version" of network cost-sharing games [Kollias and Roughgarden, 2011].

2 Preliminaries

2.1 The Shapley Value Yields Potential Games

We first review why every network cost-sharing game with Shapley cost shares is a potential game in the sense of [Monderer and Shapley, 1996]. First we define the *ordered potential* $\Phi_\sigma(S)$ with respect to an ordering σ of the players:

$$\Phi_\sigma(\mathcal{P}) = \sum_{e \in E} \sum_{\ell=1}^{|S_e|} \chi_e(i_{\sigma(\ell)}, S_{e,\sigma(1..\ell)}) \tag{3}$$

where $S_{e,\sigma(1..\ell)}$ is S_e restricted to the first ℓ in players σ. Remarkably, when the Shapley value is used as the cost-sharing method, the ordered potential (3) is the same for every ordering σ, even though individual summands generally differ [Hart and Mas-Colell, 1989, Kollias and Roughgarden, 2011]. We can therefore define $\Phi(\mathcal{P})$ as the value in (3) for an arbitrary choice of the ordering σ.

The next proposition notes that the "order-independence" property of Φ implies that it is a potential function.

Proposition 2.1 *For every network cost-sharing game with Shapley cost sharing, and for every pair $\mathcal{P}$ and $\mathcal{P}' = (P_i', P_{-i})$ that differ only in their ith component,*

$$\Phi(\mathcal{P}') - \Phi(\mathcal{P}) = c_i(\mathcal{P}') - c_i(\mathcal{P}).$$

Proof: We show that every edge e contributes the same amount to the left- and right-hand sides. If $e \in P_i \cap P_i'$ or $e \notin P_i \cap P_i'$ then we have nothing to prove. Consider $e \in P_i' \backslash P_i$. By the order-independence property of Φ_σ, we can assume without loss of generality that player i is the last player in the ordering σ, so the change in the left-hand side due to edge e is $\chi(i, S_e \cup \{i\})$, which by definition is also the change in c_i on edge e. The case for $e \in P_i \backslash P_i'$ is analogous. ■

2.2 Summability: Some Examples

Recall the definition of summability in (2).

Proposition 2.2 ([Roughgarden and Sundararajan, 2009]) *For the Shapley cost-sharing method and a set $\mathcal{C}$ of monotone cost functions, $\alpha(\mathcal{C}) \leq \mathcal{H}_k$.*

One way to prove Proposition 2.2 is to evaluate the left-hand side of (2) with a player ordering σ chosen uniformly at random. Monotonicity of C and the budget-balance property of the Shapley value imply that the ℓth summand has expected value at most $C(S)/\ell$ (conditioned on the first ℓ players, each is equally like to be last). It follows that there exists an ordering σ such that (2) holds with $\alpha = \mathcal{H}_k$, and Proposition 2.2 then follows from order-independence.

Example 2.3 (Polynomial Cost Functions) First consider anonymous polynomial cost functions, of the form $C(S) = |S|^d$ with $d \in (0, 1]$. Since the cost function is anonymous, the players all have cost shares equal to $|S|^d/|S| = |S|^{d-1}$. Since the summands are decreasing, we can upper bound the sum that ranges from 1 to $|S|$ by an integral that ranges from 0 to $|S|$:

$$\sum_{i=1}^{|S|} i^{d-1} \leq \int_0^{|S|} t^{d-1} dt$$
$$= d^{-1} t^d \Big|_0^{|S|}$$
$$= d^{-1} |S|^d,$$

which shows that $\alpha(C) \leq 1/d$, independent of k. When $d = 0$ and $C(S) = 1$ for every non-empty set S, we realize the worst case of $\alpha(C) = \mathcal{H}_k$.

More interesting is the following computation for the non-anonymous case. The summability remains $\frac{1}{d}$ with weighted players and a polynomial cost function.

Example 2.4 (Weighted Players) Suppose every player i has a weight w_i and $C(S) = \left(\sum_{i \in S} w_i\right)^d$. Shapley cost sharing yields different costs for players with different weights. However, some computations show that the summability is largest with equal-weight players (see the Appendix for a proof). Therefore, the summability with weighted players is bounded by our previous example, with $\alpha(C) \leq 1/d$.

3 Summability Characterizes Worst-Case Inefficiency

Sections 3.1–3.3 prove that, with monotone submodular cost functions, the worst-case approximation guarantee of strong Nash equilibria and potential function minimizers is exactly the summability of the Shapley cost-sharing method. Section 3.4 shows that two equilibrium refinements are incomparable on an instance-by-instance basis.

3.1 Efficiency Guarantees for Strong Nash Equilibria

The strong Price of Anarchy is the worst-case inefficiency of a strong Nash equilibrium. For network cost-sharing games, this is bounded by the summability.

Theorem 3.1 *The strong Price of Anarchy in a network cost-sharing game $\mathcal{I}_\mathcal{C}$ with submodular cost functions in $\mathcal{C}$ and Shapley cost-sharing is at most the summability $\alpha(\mathcal{C})$.*

Proof: Fix $\mathcal{I}_\mathcal{C}$, let $\mathcal{P}$ be a strong Nash equilibrium and $\mathcal{P}^*$ an optimal solution. Since $\mathcal{P}$ is a strong Nash equilibrium, for each group Γ_l of l players, there is at least one player i, who blocks changing their joint actions to $\mathcal{P}^*_{\Gamma_l}$ because $c_i(\mathcal{P}) \leq c_i(\mathcal{P}^*_{\Gamma_l}, \mathcal{P}_{-\Gamma_l})$.

Using this, we define an order σ of the players, such that $\Gamma_{i-1} = \Gamma_i \setminus \{i\}$, and player i is blocking the group Γ_i from deviating to the optimal solution.[7] Let $\mathcal{I}_{\mathcal{C},\Gamma_i}$ be a new instance in which only the members of Γ_i participate, and let $\mathcal{P}_{\Gamma_i}$ be the strategy profile $\mathcal{P}$ restricted to the players in Γ_i (which are the first i players in the order σ). The cost of player i is given by $c_i(\mathcal{P}_{\Gamma_i}) = \sum_{e \in P_i} \chi_e(i, S_{e,\Gamma_i})$, which we can also write in terms of the potential function Φ:

$$c_i(\mathcal{P}_{\Gamma_i}) = \Phi(\mathcal{P}_{\Gamma_i}) - \Phi(\mathcal{P}_{\Gamma_{i-1}}). \tag{4}$$

Here $\Phi(\mathcal{P}_{\Gamma_i})$ denotes the potential of the *restricted* game $\mathcal{I}_{\mathcal{C},\Gamma_i}$. This gives us an upper bound on the cost of player i in the original network cost-sharing game instance $\mathcal{I}_{\mathcal{C}}$ when everyone plays $\mathcal{P}$:

$$\begin{aligned}
c_i(\mathcal{P}) &\leq c_i(\mathcal{P}^*_{\Gamma_i}, \mathcal{P}_{-\Gamma_i}) && \text{(by strong Nash equilibrium)} \quad (5)\\
&\leq c_i(\mathcal{P}^*_{\Gamma_i}) && \text{(by cross-monotonicity)} \quad (6)\\
&= \Phi(\mathcal{P}^*_{\Gamma_i}) - \Phi(\mathcal{P}^*_{\Gamma_{i-1}}) && \text{(by Equation 4).} \quad (7)
\end{aligned}$$

We can now bound the total cost of the strong Nash equilibrium:

$$\begin{aligned}
C(\mathcal{P}) &= \sum_{i=1}^{k} c_i(\mathcal{P}) && \text{(by Equation 1)} \quad (8)\\
&\leq \sum_{i=1}^{k} \Big(\Phi(\mathcal{P}^*_{\Gamma_i}) - \Phi(\mathcal{P}^*_{\Gamma_{i-1}}) \Big) && \text{(by Equation 7)}\\
&= \Phi(\mathcal{P}^*_{\Gamma_k}) - \Phi(\emptyset) && \text{(by telescoping sum)}\\
&= \Phi(\mathcal{P}^*) && \text{(definition of } \Phi\text{)}\\
&= \sum_{e \in E} \sum_{i=1}^{|S_e|} \chi_e(i, S^*_{e,1..i}) && \text{(by Equation 3)}\\
&\leq \sum_{e \in E} \alpha(\mathcal{C}) \cdot C_e(S^*_e) && \text{(by summability)}\\
&= \alpha(\mathcal{C}) \cdot C(\mathcal{P}^*), && \text{(by Equation 1)}
\end{aligned}$$

and this concludes the proof. ■

3.1.1 Coalitional Smoothness

Smoothness frameworks extend price of anarchy bounds automatically to more general solution concepts. The coalitional smoothness framework of Bachrach et al. [2013] applies to equilibria that are resistant to deviations of a group of players, and the proof of Theorem 3.1 can be recast as a coalitional smoothness argument.

[7] This order is easily constructed back to front: we label a player that blocks the entire group from deviating as player k and recurse on the remaining players.

Proposition 3.2 *Network cost-sharing games with monotone submodular cost functions and Shapley cost-sharing are* $(\alpha(\mathcal{C}), 0)$*-coalitionally smooth.*

Proof: We need to show that for every ordering σ of the players, $\sum_{i=1}^{k} c_i(\mathcal{P}^*_{\Gamma_i}, \mathcal{P}_{-\Gamma_i}) \leq \alpha(\mathcal{C}) \cdot C(\mathcal{P}^*)$. While we fixed a particular order in the proof of Theorem 3.1, we only needed to do so to express c_i in terms of a deviation in (5). Since we did not use the order after this, we can conclude from the remainder of the proof that for every order, $\sum_{i=1}^{k} c_i(\mathcal{P}^*_{\Gamma_i}, \mathcal{P}_{-\Gamma_i}) \leq \alpha(\mathcal{C}) \cdot C(\mathcal{P}^*)$. ■

Since network cost-sharing games are coalitionally smooth, we inherit the extensions described in [Bachrach et al., 2013] to strong correlated equilibria [Moreno and Wooders, 1996], strong coarse correlated equilibria [Rozenfeld and Tennenholtz, 2006], and coalitional sink equilibria [Bachrach et al., 2013].

Corollary 3.3 *For network cost-sharing games:*

- *Every strong correlated equilibrium has expected cost at most* $\alpha(\mathcal{C})$ *times the optimal cost.*
- *Every strong coarse correlated equilibrium has expected cost at most* $\alpha(\mathcal{C})$ *times the optimal cost.*
- *Every coalitional sink equilibrium has expected cost at most* $\mathcal{H}_k \cdot \alpha(\mathcal{C})$ *times the optimal cost.*

We revisit the topic of strong Nash equilibria twice more in later sections. Section 4 considers cost-sharing methods other than the Shapley value, and proves that the strong Price of Anarchy remains bounded by the summability α even in games without potential functions, and Section 7 discusses conditions on the network topologies and the cost functions $\mathcal{C}$ under which strong Nash equilibria are guaranteed to exist.

3.2 Efficiency Guarantees for Potential Function Minimizers

We now consider the second equilibrium refinement, potential function minimizers. Given the existence of a potential function (Proposition 2.1), it is straightforward to apply the potential function method (see [Tardos and Wexler, 2007]) to bound the cost of its minimizers in terms of the summability. When we consider cross-monotonic cost-sharing methods other than the Shapley value in the next section, this method will no longer apply.

Theorem 3.4 *The inefficiency of the potential function minimizer in a network cost-sharing game* $\mathcal{I}_{\mathcal{C}}$ *with submodular cost functions in* $\mathcal{C}$ *and Shapley cost-sharing is at most the summability* $\alpha(\mathcal{C})$*.*

Proof: First observe the following upper bound on $\Phi(\mathcal{P})$:

$$\begin{aligned}\Phi(\mathcal{P}) &= \sum_{e \in E} \sum_{i=1}^{|S_e|} \chi_e(i, S_{e,1..i}) \\ &\leq \sum_{e \in E} \alpha(\mathcal{C}) \cdot C_e(S_e) \\ &\leq \alpha(\mathcal{C}) \cdot C(\mathcal{P}).\end{aligned}$$

Since the Shapley value is budget-balanced and cross-monotonic, $C(\mathcal{P}) \leq \Phi(\mathcal{P})$. Combining these facts and applying the potential function method, the cost of a global potential function minimizer is at most $\alpha(\mathcal{C})$ times that of an optimal solution. ■

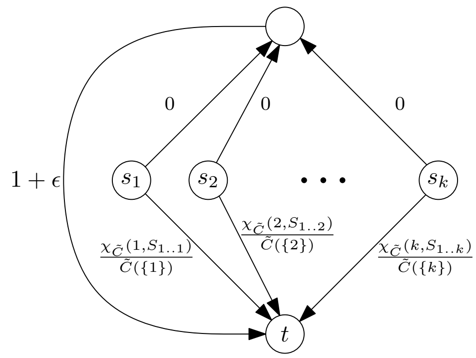

Figure 1: Tight example for Propositions 3.5 and 3.6. All cost functions are multiples of a common function $\tilde{C}$; the numbers on the edges are the multiples.

3.3 Lower Bounds

We now show how to adapt a well-known instance from [Anshelevich et al., 2008a] to show simultaneously that our two upper bounds are tight for every set $\mathcal{C}$ of monotone submodular cost functions. The key reason that the lower bound instance from [Anshelevich et al., 2008a] remains relevant in the present more general context is the cross-monotonicity of the Shapley value for a submodular cost function.

In Figure 1 we have a network with k players with different sources s_i and a common sink t. The cost functions on all edges are scalar multiples of a cost function $\tilde{C} \in \mathcal{C}$ with $\alpha(\tilde{C}) = \alpha(\mathcal{C})$ (or arbitrarily close to $\alpha(\mathcal{C})$); the numbers on the edges denote the multiples.[8]

For sufficiently small ϵ, the total cost is minimized when all players share the cost of the top route. This is not a Nash equilibrium: player k pays slightly less by taking its own personal shortcut to t. By the cross-monotonicity of the Shapley value for a submodular cost function, this holds no matter what the other players do. Given that player k takes the shortcut, player $k-1$ has a (conditional) dominant strategy to take its shortcut, and so on until player 1 does the same. This argument implies that the resulting outcome is the unique Nash equilibrium (and hence the potential function minimizer) and also a strong Nash equilibrium. The cost that each player pays is $\chi_{\tilde{C}}(i, S_{1..i})$, so by the definition of summability, the total cost of the equilibrium is, as $\epsilon \to 0$, $\alpha(\mathcal{C})$ times that the optimal cost.

Proposition 3.5 *For every set $\mathcal{C}$ of submodular cost functions and $\epsilon > 0$, there is a network cost-sharing game $\mathcal{I}_{\mathcal{C}}$ with cost functions in $\mathcal{C}$ in which the strong Price of Anarchy is at least $\alpha(\mathcal{C}) - \epsilon$.*

Proposition 3.6 *For every set $\mathcal{C}$ of submodular cost functions and $\epsilon > 0$, there is a network cost-sharing game $\mathcal{I}_{\mathcal{C}}$ with cost functions in $\mathcal{C}$ in which the cost of the potential function minimizer is at least $\alpha(\mathcal{C}) - \epsilon$ times the optimal cost.*

[8]For constructing lower bounds, we can assume without loss of generality that the set $\mathcal{C}$ is closed under scalar multiplication. The reason is that scalar multiples can be simulated by replacing edges of a lower bound example with paths of suitable lengths.

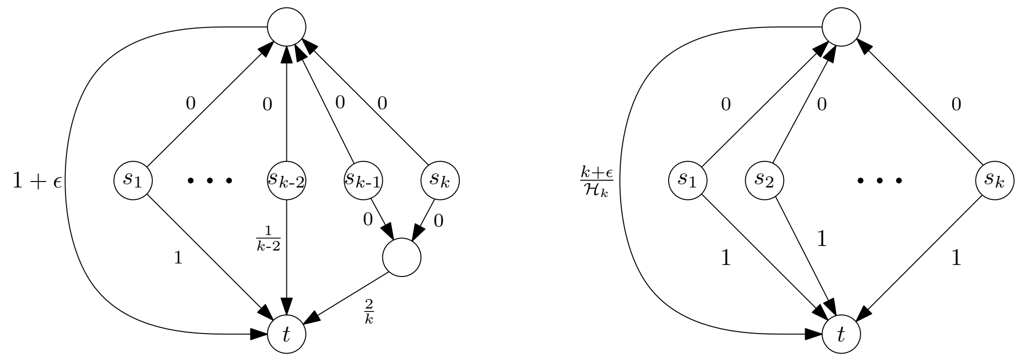

(a) The potential function minimizer is better than the Strong Price of Anarchy.

(b) The Strong Price of Anarchy is better than the potential function minimizer.

Figure 2: Instances where one equilibrium refinement yields the optimal solution, while the other exhibits worst-case inefficiency.

3.4 Comparison of Equilibrium Concepts

Although both equilibrium refinements have the same tight worst-case guarantees, the refinements are incomparable on an instance-by-instance basis. Even for network cost-sharing games with constant functions, there are instances where one refinement yields the optimal solution and the other exhibits worst-case inefficiency.

Proposition 3.7 *For all $\epsilon > 0$ and sufficiently large k, there is a network cost-sharing game with cost constant functions such that the potential function minimizer is the optimal solution and the strong price of anarchy is at least $\mathcal{H}_k - \epsilon$.*

Proof: We adapt an example from Epstein et al. [2009, Thm 3.12]; see Figure 2a. There are two Nash equilibria, one where all players take the top route $\mathcal{P}_{\text{top}}$ and one where all take the bottom route $\mathcal{P}_{\text{bottom}}$.

The strategy where all players share the cost going to t has potential $\Phi(\mathcal{P}_{\text{top}}) = (1+\epsilon)\mathcal{H}_k$, whereas the potential of all taking their own paths is $\Phi(\mathcal{P}_{\text{bottom}}) = \mathcal{H}_{k-2} + \frac{2}{k} + \frac{1}{k} > \mathcal{H}_k$ for $k > 2$. We can pick ϵ small enough so that $\Phi(\mathcal{P}_{\text{top}}) < \Phi(\mathcal{P}_{\text{bottom}})$, making $\mathcal{P}_{\text{top}}$ the potential function minimizer. However, if players k and $k-1$ deviate to their bottom strategies at the same time, they both strictly decrease their costs, hence the Nash equilibrium of the potential function minimizer is not a strong Nash equilibrium. The only strong Nash equilibrium is $\mathcal{P}_{\text{bottom}}$, which is a factor $\mathcal{H}_{k-2} + \frac{2}{k}$ worse than the optimal. ■

Proposition 3.8 *For all $\epsilon > 0$ and $k \geq 1$, there is a network cost-sharing game with cost constant functions such that every strong Nash equilibrium is an optimal solution and the potential function minimizer has cost at least $\mathcal{H}_k - \epsilon$ times that of an optimal solution.*

Proof: The network is given in Figure 2b. Again there are two Nash equilibria, $\mathcal{P}_{\text{top}}$ and $\mathcal{P}_{\text{bottom}}$. The potential of the strategy profile where every player chooses their bottom path to t is $\Phi(\mathcal{P}_{\text{bottom}}) = \sum_{i=1}^{k} 1 = k$, whereas the potential of the strategy profile where all players share the cost of the joint path to t is $\Phi(\mathcal{P}_{\text{top}}) = (\frac{k+\epsilon}{\mathcal{H}_k}) \sum_{i=1}^{k} \frac{1}{i} = k + \epsilon > \Phi(\mathcal{P}_{\text{bottom}})$. However, the potential function minimizer has total cost k, while the optimal solution has total cost $\frac{k+\epsilon}{\mathcal{H}_k}$. Since every player pays 1 on their bottom path, but only about $\frac{1}{\mathcal{H}_k}$ in the optimal solution, the optimal solution is the only strong Nash equilibrium. ■

4 Beyond Submodular Cost Functions

The previous section considered submodular cost functions and cost sharing according to the Shapley value. Recall that the Shapley value is a cross-monotonic cost-sharing method for every submodular cost function, and is not generally cross-monotonic for non-submodular cost functions. This section considers two natural directions for further generalization. Section 4.1 retains Shapley cost shares (and hence a potential function) but relaxes submodularity (and hence cross-monotonicity). Section 4.2 considers cross-monotonic (and hence non-Shapley) cost-sharing methods for non-submodular cost functions.

4.1 Non-Submodular Costs with Shapley Cost Sharing

If we use Shapley cost shares with non-submodular (monotone) cost functions, the corresponding network cost-sharing games continue to have a potential function (Proposition 2.1), the summability remains bounded by $\mathcal{H}_k$ (Proposition 2.2), and the summability continues to upper bound the worst-case approximation ratio of potential function minimizers (Theorem 3.4). However, without cross-monotonicity the lower bound in Proposition 3.6 no longer holds. The reason is that we use cross-monotonicity to argue conditional dominant strategies. The last player has a dominant strategy to take their lower path, because she knows that her most advantageous situation is when all players share the cost of the top path, and in that case she pays less by taking the bottom path. However, if her cost could decrease by another player choosing to deviate to their lower path, this is no longer a dominant strategy for her, and the equilibrium can depend on the cost function. Indeed, there are examples of monotone non-submodular cost functions such that the approximation ratio of potential function minimizes is strictly better than the summability (see [Christodoulou and Koutsoupias, 2005] for one).

Also, without cross-monotonicity, the upper bound (Theorem 3.1) for the strong price of anarchy fails to hold. The proof of Theorem 3.1 uses cross-monotonicity to upper bound the cost for a player in an entangled strategy profile in (6). Indeed, Theorem 3.1 is generally false for non-cross-monotonic cost-sharing methods (see [Chien and Sinclair, 2009, Theorem 5.1] for an example). The lower bound argument (Proposition 3.5) for the price of anarchy of strong Nash equilibria breaks down for the same reasons as for potential function minimizers.

We conclude that relaxing the cross-monotonicity constraint has significant consequences: neither the upper nor lower bounds for strong Nash equilibria carry over, and while the upper bound for potential function minimizers carries over, it is generally not tight.

4.2 Non-Shapley Cross-Monotonic Cost Sharing

This section explores non-Shapley-value cost-sharing methods that are cross-monotonic; the design of such methods has been explored extensively in the context of cost-sharing mechanisms (see Jain and Mahdian [2007] for a survey). Network cost-sharing games with arbitrary cross-monotonic cost-sharing methods no longer admit potential functions in general, and even the best PNE can be arbitrarily bad (see [Chen and Roughgarden, 2009] for an example), so we focus on strong Nash equilibria. Despite the non-existence of a potential function, we can still prove that whenever a strong Nash equilibrium exists, it is approximately optimal.

We also allow non-budget-balanced rules; this relaxation permits cross-monotonic cost-sharing rules for many non-submodular cost functions (see Jain and Mahdian [2007]).

Definition 4.1 A cost-sharing method is β-budget balanced if it recoups at least a β fraction of the cost from the players:

$$\frac{C(S)}{\beta} \leq \sum_{i \in S} \chi(i, S) \leq C(S).$$

This also means that the total cost for the game does not necessarily equal the total cost over all the players. Instead, we have

$$C(\mathcal{P}) = \sum_{e \in E} C_e(S_e) \leq \beta \sum_{i=1}^{k} c_i(\mathcal{P}).$$

Let $\mathcal{F}$ denote a class of monotone cost functions, possibly including non-submodular functions. Let $\mathcal{I}_{\mathcal{F}}$ denote an instance of a network cost-sharing game with cost functions from $\mathcal{F}$. Let $\alpha(\mathcal{F})$ denote the summability of $\mathcal{F}$. Since we no longer use the Shapley value for cost sharing, the cost-sharing method χ is no longer order-independent, and summability is defined via a worst-case order σ as in (2). The summability of a cost-sharing method other than the Shapley value can be larger than $\mathcal{H}_k$, but for many non-submodular cost functions, there exist $O(1)$-budget-balanced cross-monotonic cost-sharing methods with summability $O(\log k)$ or $O(\log^2 k)$; see [Roughgarden and Sundararajan, 2009] for a survey. Finally, we also revert back to the *ordered potential* Φ_σ in the original definition in (3).

Theorem 4.2 *For arbitrary monotone cost functions $\mathcal{F}$, and a cost-sharing method χ that is cross-monotonic and β-budget balanced, the strong Price of Anarchy in a network cost-sharing game $\mathcal{I}_{\mathcal{F}}$ with cost functions in $\mathcal{F}$ is at most $\beta \cdot \alpha(\mathcal{F})$.*

Proof: We follow the proof outline of Theorem 3.1. Fix $\mathcal{I}_{\mathcal{F}}$, let $\mathcal{P}$ be a strong Nash equilibrium and $\mathcal{P}^*$ an optimal solution. Define σ to be an order such that for every ℓ, $\sigma(\ell)$ blocks the joint deviation of $\Gamma_{\sigma(\ell)}$ to $\mathcal{P}^*$. The ordered potential function Φ_σ can be used to bound the cost of the *last* player in σ:

$$c_{\sigma(\ell)}(\mathcal{P}_{\Gamma_{\sigma(\ell)}}) = \Phi_\sigma(\mathcal{P}_{\Gamma_{\sigma(\ell)}}) - \Phi_\sigma(\mathcal{P}_{\Gamma_{\sigma(\ell-1)}}).$$

Player $\sigma(\ell)$ is unique in this, and the identity does not hold for the players preceding it. We can therefore bound the cost of player $\sigma(\ell)$ using the ordered potential:

$$c_{\sigma(\ell)}(\mathcal{P}) \leq c_{\sigma(\ell)}(\mathcal{P}^*_{\Gamma_{\sigma(\ell)}}, \mathcal{P}_{-\Gamma_{\sigma(\ell)}}) \leq c_{\sigma(\ell)}(\mathcal{P}^*_{\Gamma_{\sigma(\ell)}}) = \Phi_\sigma(\mathcal{P}^*_{\Gamma_{\sigma(\ell)}}) - \Phi_\sigma(\mathcal{P}^*_{\Gamma_{\sigma(\ell-1)}}).$$

Next, following (8) but using that χ is β-budget balanced, we obtain

$$C(\mathcal{P}) \leq \beta \sum_{\ell=1}^{k} c_{\sigma(\ell)}(\mathcal{P}) = \beta \sum_{i=1}^{k} \Big(\Phi_\sigma(\mathcal{P}^*_{\Gamma_i}) - \Phi_\sigma(\mathcal{P}^*_{\Gamma_{i-1}}) \Big) = \beta \Phi_\sigma(\mathcal{P}^*) \leq \beta \cdot \alpha(\mathcal{F}) \cdot C(\mathcal{P}^*),$$

where the last inequality holds because $\alpha(\mathcal{F})$ is defined as the supremum over all orders σ. ■

Proposition 3.2 and Corollary 3.3 also carry over to the present more general setting, with a loss of β in the approximation factors.

5 Negative Results for Non-Anonymous Cost Functions

This section shows that, in contrast to the results of Section 3, not all positive results known for network cost-sharing games with anonymous submodular cost functions carry over to the non-anonymous submodular case. We consider both the greedy algorithm of Syrgkanis [2010] for computing a low-cost PNE, and the equilibrium refinement implicit in the subgame perfect equilibria identified by Paes Leme et al. [2012]. Section 5.1 explains the models and results from previous work. Section 5.2 shows that these results do not extend to non-anonymous submodular cost functions.

5.1 Model and Previous Work

5.1.1 Singleton Cost-Sharing Games

This section restricts attention to *singleton* cost-sharing games, where each strategy of each player is a single edge. Different players can have different strategy sets. Such games can be thought of as k players that each have a task that they need to run on one of m machines, with a bipartite graph indicating which tasks can run on which machines. The cost of a machine j is shared by all players that use that machine. It is straightforward to model such a game as a network cost-sharing game.

5.1.2 A Greedy Algorithm

Syrgkanis [2010] gave a greedy algorithm for computing a PNE of such a game that has cost bounded above by the potential function value of an optimal solution. This algorithm proceeds analogously to the greedy algorithm for the set cover problem, where sets correspond to machines and elements to players. Each iteration, the algorithm picks one machine and assigns as many unassigned players to it as possible; subject to this, it picks the machine that minimizes the cost share of the newly assigned players. We denote this algorithm by Greedy .

Theorem 5.1 (Syrgkanis [2010]) *In singleton anonymous cost-sharing games, the* Greedy *algorithm computes a PNE with cost at most the potential function value of an optimal solution.*

As in the proof of Theorem 3.4, this translates to an upper bound in terms of the summability of the cost functions in $\mathcal{C}$.

Corollary 5.2 *In singleton anonymous cost-sharing games with cost functions in $\mathcal{C}$, the* Greedy *algorithm computes a PNE with cost at most $\alpha(\mathcal{C})$ times that of an optimal solution.*

Since the lower bound instance in Propositions 3.5 and 3.6 can be realized by a singleton cost-sharing game, Corollary 5.2 is tight for every set $\mathcal{C}$ of monotone anonymous submodular cost functions.

5.1.3 Sequential Price of Anarchy

We next describe a sequential version of a singleton cost-sharing game. There is an exogenous ordering on the players. Player i can observe all the actions $a_j \in A_j$ of players before her, and the utility functions $u_i(a_1, ..., a_k)$ of all players i is public knowledge. A strategy S_i is a mapping from the actions of preceding players $a_1, ..., a_{i-1}$ to an action $a_i \in A_i$.

Given a prefix $(\alpha_{1..j-1}) \in A_1 \times .. \times A_{j-1}$, the induced subgame for players $j, ..., k$ is given by strategies $S_i(\alpha_1, .., \alpha_{j-1}, a_j, .., a_{i-1})$ and utility functions $u_i(\alpha_1, .., \alpha_{j-1}, a_j, .., a_k)$. A subgame perfect equilibrium is a strategy profile S that simultaneously is an equilibrium for all of its subgames. The sequential price of anarchy is the ratio of the worst subgame perfect equilibrium and the optimal solution.

Previous results of Syrgkanis [2010] and Paes Leme et al. [2012] imply that summability bounds from above the sequential price of anarchy in singleton anonymous cost-sharing games.[9]

Proposition 5.3 *For singleton cost-sharing games with anonymous cost functions in $\mathcal{C}$, the sequential price of anarchy is at most $\alpha(\mathcal{C})$.*

Proof: Paes Leme et al. [2012, Theorem 1] prove that, for every player ordering, the Greedy algorithm yields the unique subgame perfect equilibrium, with the ith player playing the strategy computed for it by the algorithm. The proposition follows from Corollary 5.2. ■

Adapting the standard lower bound example of Figure 1 shows that the bound in Proposition 5.3 is tight.

The proof of Proposition 5.3 shows that, ranging over all possible player orderings in an anonymous singleton cost-sharing game, all subgame perfect equilibria induce the same single-shot outcome, which is a PNE. In this sense, it can be interpreted as an upper bound of $\alpha(\mathcal{C})$ on an equilibrium refinement in the single-shot game — the PNE induced by subgame perfect equilibria or, equivalently, the PNE computed by the Greedy algorithm.

5.2 Non-anonymous Cost Functions

Unfortunately, the results surveyed in Section 5.1 do not extend to singleton cost-sharing games with non-anonymous submodular cost functions.

5.2.1 The Greedy Algorithm

There are multiple ways to extend the greedy algorithm to singleton cost-sharing games with non-anonymous cost functions and Shapley cost shares. The reason is that when many players are assigned to a common machine in a single iteration, they generally incur different cost shares. The next proposition notes that the three arguably most natural extensions — minimizing the minimum, the maximum, or the average cost share of a newly assigned player — do not yield algorithms guaranteed to output a PNE.

Proposition 5.4 *For a greedy algorithm that uses the "lowest average cost," "lowest minimum cost," or "lowest maximum cost" rules, there exists a singleton cost-sharing game with submodular cost functions and Shapley cost sharing such that the greedy algorithm does not yield a PNE.*

[9] This result requires a generalized "no ties" assumption on the cost functions; see Paes Leme et al. [2012] and Bilò et al. [2013] for details.

Proof: We give two examples: the first shows a singleton cost-sharing game where the "lowest average cost" and "lowest maximum cost" rules make the wrong decision, the second example shows that the "lowest minimum cost" rule can also make the wrong decision. Both examples have three players and two machines: player p_1 can only choose machine a, player p_2 can only choose machine b, and player p_3 can choose either. Both machines have the same cost function.

Consider the following submodular cost function:

$$
\begin{aligned}
C(\{p_1\}) &= 1 & C(\{p_1,p_2\}) &= 1 & & \\
C(\{p_2\}) &= 3/4 & C(\{p_1,p_3\}) &= 1+\epsilon & C(\{p_1,p_2,p_3\}) &= 1+\epsilon \\
C(\{p_3\}) &= 1/2 & C(\{p_2,p_3\}) &= 1. & &
\end{aligned}
$$

Both the "lowest average cost" and "lowest maximum cost" rules assign p_3 to machine b (with p_2), but this is not a Nash equilibrium.

Similarly for the "lowest minimum cost" rule, for the submodular cost function

$$
\begin{aligned}
C(\{p_1\}) &= 1/2 & C(\{p_1,p_2\}) &= 1 & & \\
C(\{p_2\}) &= 3/4 & C(\{p_1,p_3\}) &= 1 & C(\{p_1,p_2,p_3\}) &= 1 \\
C(\{p_3\}) &= 1 & C(\{p_2,p_3\}) &= 1, & &
\end{aligned}
$$

player p_3 is assigned to machine a with p_1, while p_3 will want to switch to machine b. ■

5.2.2 Subgame Perfect Equilibria

With non-anonymous submodular cost functions and Shapley cost shares, it is no longer true that all subgame perfect equilibria (across all player orderings) of a singleton cost-sharing game induce a common PNE, or even a PNE at all.

Proposition 5.5 *Different player orderings can lead to different subgame perfect equilibria in singleton cost-sharing games with non-anonymous submodular cost functions and Shapley cost shares.*

Proof: Consider two players and two machines, a and b. The first machine has cost function

$$
\begin{aligned}
C_a(\{p_1\}) &= 4/6 \\
C_a(\{p_2\}) &= 5/6 \\
C_a(\{p_1,p_2\}) &= 1,
\end{aligned}
$$

and machine b has the roles of p_1 and p_2 reversed. Regardless of who goes first, both players are better off by sticking together. However, p_1 would prefer to be on machine a, whereas p_2 wants to be on machine b. Therefore whoever goes first in the order determines which equilibrium the players converge to. ■

Lemma 5.6 *In a singleton cost-sharing game with non-anonymous submodular cost functions and Shapley cost shares, the single-shot outcome induced by a subgame perfect equilibrium may not be a PNE.*

Proof: Consider three machines a, b, c and three players with strategy spaces $\{a, b\}$, $\{a, c\}$, and $\{b, c\}$. Machines a and c have anonymous cost functions, and machine b has a non-anonymous one:

$$
\begin{array}{lll}
C_a(1) = 3/2 + \epsilon & C_b(\{p_1\}) = 3/2 + 2\epsilon & C_c(1) = 3/2 \\
C_a(2) = 2 & C_b(\{p_3\}) = 2 & C_c(2) = 2 + 2\epsilon \\
& C_b(\{p_1, p_3\}) = 2. &
\end{array}
$$

Take the order p_2, p_1, p_3. Solving for the subgame perfect equilibrium yields p_2 on c, p_1 on a, and p_3 on c. This is not a Nash equilibrium as p_2 would prefer to switch to a. ■

We conclude that the positive results known for singleton cost-sharing games with anonymous submodular cost functions cannot be extended, without significant compromise, to non-anonymous submodular cost functions with Shapley cost shares.

6 Equivalence between Network Cost-Sharing Games and Moulin Mechanisms

The concept of summability was first proposed in the context of cost-sharing mechanisms [Roughgarden and Sundararajan, 2009]. This section makes explicit a connection between Moulin mechanisms [Moulin, 1999, Moulin and Shenker, 2001] and network cost-sharing games. We first briefly discuss cost-sharing mechanisms and then show a sense in which Moulin mechanisms are special cases of network cost-sharing games. In particular, negative results for the inefficiency of Moulin mechanisms transfer directly to network cost-sharing games.

6.1 Moulin Mechanisms

There are n players, and each has a valuation v_i for "winning" (and 0 for "losing"). A cost function C indicates the cost incurred as a function of the set S of winners. A Moulin mechanism is defined by a cross-monotonic cost-sharing method $\chi(i, S)$. An outcome is then the choice S of a winning set, with player $i \in S$ paying $\chi(i, S)$ (and losers paying nothing).

Moulin [1999] described two closely related models. In the *full-information model*, we treat players' valuations as common knowledge. Each player has two actions, "in" and "out." A player's payoff is 0 when it chooses "out," and is $v_i - \chi(i, S)$ when it chooses "in," where S denotes the set of players choosing "in." Cross-monotonicity implies that the union of PNE is again a PNE, so there is a unique maximal PNE in such a game. This maximal PNE is also strong.

In the *private-information model*, we treat players' valuations as private. The strategy of a player is now a nonnegative bid. A Moulin mechanism takes these bids at face value and computes the corresponding maximal strong PNE. This mechanism is strategyproof, meaning that for every player, it is a dominant strategy to bid its true valuation. The mechanism is also group-strategyproof, meaning that coalitions that do not use side payments also have no incentive to collectively misreport valuations.

Following Roughgarden and Sundararajan [2009], we compare different Moulin mechanism outcomes with the social cost objective function:

$$\pi(S) = C(S) + \sum_{i \notin S} v_i,$$

where S denotes the set of "in" players.

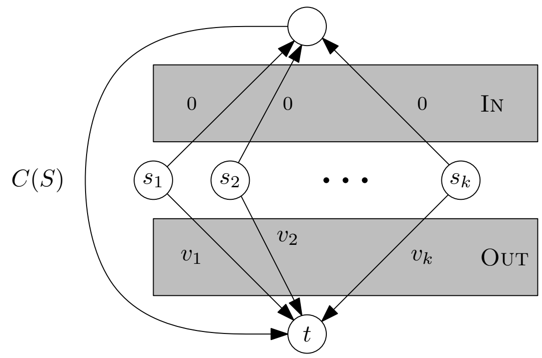

Figure 3: The network cost-sharing game that is equivalent to any in/out cost-sharing game.

6.2 Equivalence

By construction, there is a natural equivalence between the full-information and private-information models of Moulin mechanisms. Next, we show an equivalence between the full-information model and the subclass of network cost-sharing games that are played on the network of Figure 3. By "equivalent," we mean there is a bijection between players and strategy sets that preserves the objective function value and players' costs (or negative payoffs), up to a universal constant.

Proposition 6.1 *Full-information Moulin mechanisms are equivalent to a subclass of network cost-sharing games.*

Proof: Figure 3 depicts the equivalence. The direct path s_i to t corresponds to the "out" strategy of player i at opportunity cost v_i. The other strategy corresponds to the "in" strategy. There is an obvious bijection between players and strategy sets. The social cost of the Moulin mechanism outcome equals the cost of the corresponding network cost-sharing game outcome. Up to an additive constant (of v_i), this bijection preserves the costs incurred by all players. ■

Proposition 6.1 implies that lower bounds for the social cost approximation of Moulin mechanism outcomes apply equally well to network cost-sharing games.

Corollary 6.2 *Let C be a monotone cost function and χ a cost-sharing method. If the worst-case social cost approximation of the corresponding Moulin mechanism is at least ρ, then the worst-case strong price of anarchy in network cost-sharing games with cost function C and cost-sharing method χ is at least ρ.*

Since summability characterizes the worst-case social cost approximation of Moulin mechanisms [Roughgarden and Sundararajan, 2009], the lower bound of Corollary 6.2 matches that of Proposition 3.5. Put differently, the subclass of network cost-sharing games corresponding to Moulin mechanisms is guaranteed to furnish worst-case examples for the strong price of anarchy (amongst all network cost-sharing games).

7 Existence of Strong Equilibria

Sections 3.1 and 4 showed that when strong Nash equilibria exist, their cost is bounded by the summability $\alpha(\mathcal{C})$ (respectively, $\alpha(\mathcal{F})$) times that of an optimal outcome. One of the criticisms of strong Nash equilibria is that they are not guaranteed to exist in every instance. Epstein et al. [2009] showed that for constant cost functions with equal cost-sharing, strong Nash equilibria are guaranteed to exist in all graphs when all players have the same source and sink; in series-parallel graphs when all players share the same source; and in extension-parallel graphs for arbitrary sources and sinks for all players.

In this section we show that the generalization from equal cost-sharing of constant functions to anonymous cross-monotonic cost sharing yields guaranteed existence for the same network topologies. However, once we allow non-anonymous cost sharing, existence is no longer guaranteed for these topologies. We give an example of an extension-parallel graph with 2 players that share the same source and sink, which admits no strong Nash equilibrium.

For the existence questions studied in this section, only the cost-sharing method χ is relevant; players care about the cost functions only through their cost shares. The cost shares we use are not pathological: all of them arise as Shapley values with respect to a submodular cost function.

7.1 Anonymous Cost-Sharing Methods

In the case where we have different cost-sharing methods on the edges, but each cost-sharing method shares the cost equally among all players, we can extend the proofs from Epstein et al. [2009] quite straightforwardly. We give the proof for the single-source single-sink case directly; for single-source and multi-sink networks we give the lemma that allows Epstein et al.'s proofs to go through.

Lemma 7.1 *For a network cost-sharing game with possibly different anonymous cross-monotonic methods, where all players have the same source and sink, there exists a strong Nash equilibrium in which all players take the same path.*

Proof: Take the entire set of players, and pick the path P that has the lowest total cost for the entire group. This is a strong Nash equilibrium. For if not, there is a deviation to one or more paths where all players in the coalition decrease their cost. Take a deviating path P' that decreased the cost for a player i. By anonymity and cross-monotonicity of the cost-sharing method, any other player can join this path for at most the cost of player i. Therefore, all the players can move to this path, to decrease their cost. Hence the path P could not have been the minimizing path. ■

The remainder of the proofs in [Epstein et al., 2009] use the recursive definition of series and extension parallel graphs. Their proof for series composition relies on a key lemma that shows that players whose source and sink are the source and sink of the graph minimize their cost in one of the strong Nash equilibria. We extend their proof from equal sharing of constant cost functions, to anonymous cross-monotonic cost-sharing methods.

Lemma 7.2 *Consider a series-parallel graph with possibly different anonymous cross-monotonic methods, source s, and sink t. Suppose player i has source $s_i = s$ and sink $t_i = t$ and that the instance has at least one strong Nash equilibrium. Then, player i's minimum cost across all strong Nash equilibria equals its minimum cost across all strategy profiles.*

Proof: Let $\mathcal{P}$ denote a strong Nash equilibrium that minimizes the cost of player i. Let $\bar{\mathcal{P}}$ denote a strategy profile that minimizes the cost of player i. Among all such strategy profiles, choose that minimizes the number of players that choose different strategies in $\mathcal{P}$ and $\bar{\mathcal{P}}$; let Γ denote these players. If there is such a strategy profile with $i \in \Gamma$, then choose that strategy profile.

The proof works as follows: we create a reference profile $\mathcal{P}^*$ where player i plays $\bar{P}_i$ and all players in Γ interact as much as possible with player i. (Players outside of $\Gamma \cup \{i\}$ play as in $\mathcal{P}$ and $\bar{\mathcal{P}}$.) By [Epstein et al., 2009, Lemma A.1], the first and last vertex where players can interact with i are uniquely determined in a series-parallel graph: players align their path between these two vertices, and cannot interact any earlier or later with the path of player i. Due to cross-monotonicity of the cost-sharing methods, $c_i(\mathcal{P}^*) \leq c_i(\bar{\mathcal{P}})$. Since i's cost in $\bar{\mathcal{P}}$ is minimal over all profiles, so is its cost in $\mathcal{P}^*$.

Since $\mathcal{P}$ is a strong Nash equilibrium, there is a player $j \in \Gamma$ with $c_j(\mathcal{P}^*) \geq c_j(\mathcal{P})$. We will show that we can either remove j from Γ, or swap her for i to arrive at a contradiction. Let $\mathcal{P}'$ be the strategy profile where j plays P_j and all other players in Γ interact maximally with j. By cross-monotonicity, we have $c_j(\mathcal{P}') \leq c_j(\mathcal{P}) \leq c_j(\mathcal{P}^*)$. Note that j plays P_j, and only players from Γ changed their profile with respect to $\mathcal{P}$. If $i \in \Gamma$, then for the subpaths where i and j interact $c_i(\mathcal{P}') = c_j(\mathcal{P}') \leq c_j(\mathcal{P}) \leq c_j(\mathcal{P}^*) = c_i(\mathcal{P}^*)$, where the equalities hold because the cost-sharing methods are anonymous. Therefore $\mathcal{P}'$ also minimizes c_i with a smaller group Γ, contradicting our choice of $\bar{\mathcal{P}}$. If $i \neq \Gamma$, then we can put i in Γ on the path of j, where again we get for the subpaths where they interact $c_i(\mathcal{P}') = c_j(\mathcal{P}') \leq c_j(\mathcal{P}) \leq c_j(\mathcal{P}^*) = c_i(\mathcal{P}^*)$, reaching the same contradiction. ■

Combining Lemma 7.2 with the arguments in Epstein et al. [2009] yields the following.

Proposition 7.3 *For every network cost-sharing game on a series-parallel graph with a common source vertex and with possibly different anonymous cross-monotonic cost-sharing methods, there is at least one strong Nash equilibrium.*

Proposition 7.4 *For every network cost-sharing game on an extension-parallel graph with possibly different anonymous cross-monotonic cost-sharing methods, there is at least one strong Nash equilibrium.*

7.2 Non-anonymous Cost Sharing

With non-anonymous cost sharing methods, we lose guaranteed strong Nash equilibria even in very restricted settings.

Lemma 7.5 *There exists an extension-parallel graph with submodular cost functions and Shapley cost shares, and two players with a common source and sink, such that no strong Nash equilibrium exists.*

Proof: We adapt the example from Rozenfeld and Tennenholtz [2006], see Figure 4. There are two players p_1, p_2 on the graph. There is only one Nash equilibrium, which is that p_1 takes edge e_1, and p_2 takes edges e_2 and e_3; both for a cost of 4. If they both deviate to edges e_2, e_4, then their cost would decrease to 3, and hence there is no strong Nash equilibrium. ■

On the positive side, an easy argument shows that if all players have a common source and sink, and all edges have the same cross-monotonic cost-sharing method, then strong Nash equilibria exist.

Proposition 7.6 *For every network cost-sharing game with a common source and sink vertex and a common (possibly non-anonymous) cross-monotonic cost-sharing method, there is at least one strong Nash equilibrium. In particular, it is a strong Nash equilibrium for all players to use a common source-sink path with the minimum-possible number of edges.*

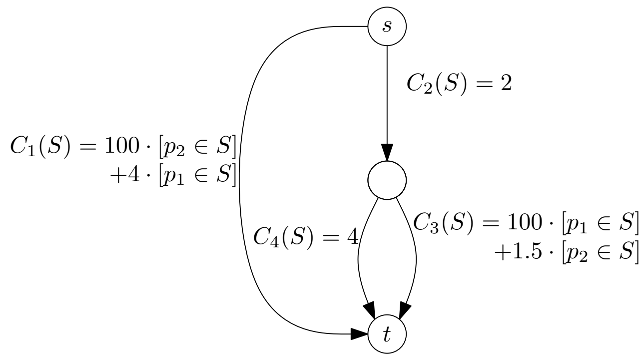

Figure 4: There is no strong Nash equilibrium on this network. The only Nash equilibrium is for p_1 to take e_1, and p_2 to take edges e_2 and e_3, but both can decrease their cost by sharing edges e_2 and e_4.

8 Conclusions

This paper studied systematically network cost-sharing games with non-anonymous cost functions; the large literature on network cost-sharing games has confined its attention almost entirely to the anonymous case. Non-anonymous cost functions arise naturally when, for example, different players have different sizes or service requirements.

Two well-studied equilibrium refinements, strong Nash equilibria and potential function minimizers, are near-optimal with non-anonymous cost functions as long as the functions are submodular and costs are shared using the Shapley value. Strong Nash equilibria remain near-optimal, when they exist, provided costs are shared by a cost-monotonic cost-sharing method. All of these worst-case approximation guarantees equal the summability of the cost-sharing method used for the set of allowable cost functions. This exact characterization unifies our understanding of these equilibrium refinements in network cost-sharing games with the theory of efficiency guarantees for Moulin mechanisms.

A Summability for Weighted Players and Polynomial Cost

This Appendix supplies the computations needed to justify the assertion in Example 2.4 that, for a fixed polynomial cost function, the summability is maximized by equal-weight players.

First, we recall that Kollias and Roughgarden [2011] show that the sum of Shapley cost shares can be written as a linear combination of the costs of subsets. The coefficient of a subset of size $|T|$ is $a_T = \frac{(|T|-1)!}{k(k-1)\cdots(k-|T|+1)}$, which we can rewrite as follows:

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

This gives us an alternative formulation for the left hand side of the summability definition (2) when χ is the Shapley value:

$$\sum_{i=1}^{k} \chi(i, S_i) = \sum_{i=1}^{k} \frac{1}{i} \frac{1}{\binom{k}{i}} \sum_{T \subseteq S, |T|=i} C(T).$$

Consequently, the summability is given by

$$\alpha(\mathcal{C}) = \sup_S \frac{1}{C(S)} \sum_{i=1}^{k} \frac{1}{i} \frac{1}{\binom{k}{i}} \sum_{T \subseteq S, |T|=i} C(T). \tag{9}$$

It is straightforward to show that, in the case of weighted players and a polynomial cost function, scaling the weights has no effect on the summability.

Proposition A.1 *The summability $\alpha(\mathcal{C})$ of $C(S) = \left(\sum_{i\in S} w_i\right)^d$ with Shapley cost sharing is invariant under uniform scaling of the weights w_i.*

We can now prove that over all weight profiles, uniform weights maximize the summability.

Lemma A.2 *The summability $\alpha(\mathcal{C})$ of $C(S) = \left(\sum_{i\in S} w_i\right)^d$ and Shapley cost sharing is maximized when all weights are equal.*

Proof: Order the weights so that $w_1 \geq w_2 \geq .. \geq w_k$. By Proposition A.1, we can assume without loss of generality that $\sum_i w_i = 1$. Assume that a weight profile with $w_1 > w_k$ yields the maximum summability.[10] Rewrite $w_1 = w+\epsilon, w_k = w-\epsilon$; we show that the derivative of the summability with respect to ϵ is strictly negative for all sufficiently small $\epsilon > 0$, contradicting that the summability was maximized.

Since we normalized the weights, the denominator in Equation (9) is 1, so the summability is given by $\sum_{i=1}^{k} \frac{1}{i}\frac{1}{\binom{k}{i}} \sum_{T\subseteq S,|T|=i} \left(\sum_{i\in T} w_i\right)^d$. For every summand $C(T)$ that includes both or none of player 1 and player k, we have $\frac{d}{d\epsilon}C(T) = 0$. For every T containing neither 1 nor k, the summands $C(T\cup\{1\})$ and $C(T\cup\{k\})$ have the same coefficient. We pair these up with each other and show that $\frac{\mathrm{d}}{\mathrm{d}\epsilon}C(T\cup\{1\}) + C(T\cup\{k\}) < 0$ for $\epsilon > 0$. Recall that $w_1 = w+\epsilon, w_k = w-\epsilon$ and let $w' = w + \sum_{i\in T} w_i$. Then,

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\epsilon}C(T\cup\{w_1\}) + C(T\cup\{w_k\}) &= \frac{\mathrm{d}}{\mathrm{d}\epsilon}\left[(w'+\epsilon)^d + (w'-\epsilon)^d\right] \\
&= d(w'+\epsilon)^{d-1} - d(w'-\epsilon)^{d-1} \\
&< 0
\end{aligned}$$

where the last step follows because for $\epsilon > 0$, and the exponent $d-1 < 0$.[11] It follows that the first summand is smaller than the second summand, which completes the contradiction. ■

Corollary A.3 *The summability $\alpha(\mathcal{C})$ of $C(S) = \left(\sum_{i\in S} w_i\right)^d$ and Shapley cost sharing is at most $1/d$.*

[10] This maximum is attained because we can restrict attention to weight vectors that lie in the (compact) simplex and the summability is a continuous function of the weight vector.

[11] If $d = 1$ the cost share of player i is always w_i since the increase in cost function is always w_i. Therefore the summability $\alpha(\mathcal{C}) = 1$ always, including when all weights are equal.

References

E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4): 1602–1623, 2008a.

E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008b.

A. Asadpour and A. Saberi. On the inefficiency ratio of stable equilibria in congestion games. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*, pages 545–552, 2009.

R.J. Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games*, 4:287–324, 1959.

Y. Bachrach, V. Syrgkanis, E. Tardos, and M. Vojnovic. Strong price of anarchy and coalitional dynamics. arXiv:1307.2537, 2013.

V. Bilò and R. Bove. Bounds on the price of stability of undirected network design games with three players. *Journal of Interconnection Networks*, 12(1-2):1–17, 2011.

V Bilò, I. Caragiannis, A. Fanelli, and G. Monaco. Improved lower bounds on the price of stability of undirected network design games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 90–101, 2010.

V. Bilò, M. Flammini, G. Monaco, and L. Moscardelli. Some anomalies of farsighted strategic behavior. In *Proceedings of the 11th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 229–241, 2013.

L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5 (3):387–424, 1993.

M. Charikar, H. Karloff, C. Mathieu, J. Naor, and M. E. Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 70–76, 2008.

C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.

H. Chen and T. Roughgarden. Network design with weighted players. *Theory of Computing Systems*, 45(2):302–324, 2009.

H. Chen, T. Roughgarden, and G. Valiant. Designing network protocols for good equilibria. *SIAM Journal on Computing*, 39(5):1799–1832, 2010.

R. Chen and Y. Chen. The potential of social identity for equilibrium selection. *American Economic Review*, 101(6):2562–2589, 2011.

S. Chien and A. Sinclair. Strong and Pareto price of anarchy in congestion games. In *Proceedings of the 36th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 279–291, 2009.

G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, pages 59–70, 2005.

G. Christodoulou, C. Chung, K. Ligett, E. Pyrga, and R. Van Stee. On the price of stability for undirected network design. In *Proceedings of the 8th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 86–97, 2010.

A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, 67(1):51–68, 2009.

A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. J. Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *Proceedings of the 33rd International Colloqium on Automata, Languages and Programming (ICALP)*, pages 608–618, 2006.

R. Gopalakrishnan, J. R. Marden, and A. Wierman. Potential games are necessary to ensure pure Nash equilibria in cost sharing games. *Mathematics of Operations Research*, 2014. To appear.

S. Hart and A. Mas-Colell. Potential, value, and consistency. *Econometrica*, 57(3):589–614, 1989.

R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21(1):85–101, 1997.

M. O. Jackson. *Social and Economic Networks*. Princeton, 2008.

K. Jain and M. Mahdian. Cost sharing. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 15, pages 385–410. Cambridge, 2007.

E. Kalai and D. Samet. On weighted Shapley values. *International Journal of Game Theory*, 16 (3):205–222, 1987.

Yasushi Kawase and Kazuhisa Makino. Nash equilibria with minimum potential in undirected broadcast games. *Theoretical Computer Science*, 482:33–47, 2013.

K. Kollias and T. Roughgarden. Restoring pure equilibria to weighted congestion games. In *Proceedings of the 38th Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 539–551, 2011.

E. Lee and K. Ligett. Improved bounds on the price of stability in network cost sharing games. In *Proceedings of the 14th ACM conference on Electronic Commerce*, pages 607–620, 2013.

J. Li. An upper bound on the price of stability for undirected Shapley network design games. *Information Processing Letters*, 109(15):876–878, 2009.

D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

D. Moreno and J. Wooders. Coalition-proof equilibrium. *Games and Economic Behavior*, 17(1): 80–112, 1996.

H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.

H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

R. Paes Leme, V. Syrgkanis, and E. Tardos. The curse of simultaneity. In *Proceedings of the 3rd Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 60–67, 2012.

T. Roughgarden and M. Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM*, 56(4), 2009. Article 23.

O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE)*, pages 74–86, 2006.

L.S. Shapley. *Additive and Non-Additive Set Functions*. PhD thesis, Department of Mathematics, Princeton University, 1953.

M. E. Slade. What does an oligopoly maximize? *Journal of Industrial Economics*, 42:45–61, 1994.

V. Syrgkanis. The complexity of equilibria in cost sharing games. In *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE)*, pages 366–377, 2010.

É. Tardos and T. Wexler. Network formation games and the potential function method. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 19, pages 487–516. Cambridge, 2007.

P. von Falkenhausen and T. Harks. Optimal cost sharing for resource selection games. *Mathematics of Operations Research*, 38(1):184–208, 2013.