

Cost Sharing Mechanisms for Fair Pricing of Resource Usage*

Marios Mavronicolas[†] Panagiota N. Panagopoulou^{‡§} Paul G. Spirakis^{‡§}

*A preliminary version of this work appeared in the *Proceedings of the 1st International Workshop on Internet and Network Economics*, X. Deng and Y. Ye eds., pp. 210–224, Vol. 3828, Lecture Notes in Computer Science, Springer-Verlag, December 2005. This work has been partially supported by the EU within the 6th Framework Programme under contract 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS), by the General Secretariat for Research and Technology of the Greek Ministry of Development within the programme PENED 2003, and by research funds at University of Cyprus.

[†]Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus. Currently visiting Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, 33102 Paderborn, Germany. Email: mavronic@ucy.ac.cy

[‡]Department of Computer Engineering and Informatics, Patras University, Rion, Patras GR 26500, Greece.

[§]Research Academic Computer Technology Institute, Patras University, Rion, Patras GR 26500, Greece. Email: {panagopp,spirakis}@cti.gr

Abstract

We propose a simple and intuitive *cost mechanism* which assigns *costs* for the competitive usage of m *resources* by n selfish *agents*. Each agent has an individual *demand*; demands are drawn according to some probability distribution. The cost paid by an agent for a resource it chooses is the total demand put on the resource divided by the number of agents who chose that same resource. So, resources charge costs in an equitable, fair way, while each resource makes no *profit* out of the agents.

We call our model the *Fair Pricing* model. Its fair cost mechanism induces a non-cooperative game among the agents. To evaluate the *Nash equilibria* of this game, we introduce the *Diffuse Price of Anarchy*, as an extension of the *Price of Anarchy* that takes into account the probability distribution on the demands. We prove:

- *Pure Nash equilibria* may not exist, unless all chosen demands are identical.
- A *fully mixed Nash equilibrium* exists for all possible choices of the demands. Further on, the fully mixed Nash equilibrium is the *unique* Nash equilibrium in case there are only two agents.
- In the *worst-case* choice of demands, the *Price of Anarchy* is $\Theta(n)$; for the special case of two agents, the Price of Anarchy is less than $2 - \frac{1}{m}$.
- Assume now that demands are drawn from a *bounded, independent probability distribution*, where all demands are *identically distributed*, and each demand may not exceed some (*universal* for the class) constant times its expectation. It happens that the constant is just 2 when each demand is distributed symmetrically around its expectation. We prove that, for asymptotically large games where the number of agents tends to infinity, the Diffuse Price of Anarchy is at most that universal constant. This implies the *first* separation between Price of Anarchy and Diffuse Price of Anarchy.

Towards the end, we consider two closely related cost sharing models, namely the *Average Cost Pricing* and the *Serial Cost Sharing* models, inspired by Economic Theory. In contrast to the Fair Pricing model, we prove that pure Nash equilibria do exist for both these models.

1 Introduction

1.1 Motivation, Framework and Overview

We propose and analyze a very simple and intuitive *cost mechanism* for pricing the competitive usage of a collection of m *resources* by a collection of n selfish *agents*, each coming with an individual *demand*. Demands are drawn according to some (perhaps unknown) probability distribution. This assumption is suitable for many practical situations (e.g., selfish sharing of bandwidth) where repeatedly revealing the demands allows their statistical modelling.

The key feature of our mechanism is its reliance on a very natural *fairness* principle in *cost sharing*. Namely, the cost charged to an agent for a resource it chooses is the total demand on the resource divided by the number of agents who chose that same resource; we call it *Resource Cost*.

Such a cost mechanism represents a very natural sharing scheme that is often used in real life; for example, restaurants offering an “all-you-can-eat” buffet charge a fixed price to each customer, which is calculated to alleviate all restaurant costs over customers. Moreover, Internet service providers and operators in telecommunication networks often charge a flat amount in case the demands of agents on bandwidth do not differ much – see, for example, [21, 22], or the more recent [20] and references therein. Moreover, this cost mechanism represents a *fair* pricing scheme since no resource makes any profit by overcharging, while all agents sharing the same resource are treated equally.*

In reflection to the fairness principle encapsulated in its cost mechanism, our pricing model will be coined as the *Fair Pricing* model. Its fair cost mechanism induces a non-cooperative strategic game, which we call *FairPricingGame*, whose *players* and *strategies* are the agents and resources, respectively. We analyze the *Nash equilibria* [18] (both *pure* and *mixed*) for *FairPricingGame*; roughly speaking, these are stable states from which no agent has an incentive to unilaterally deviate. In particular, we consider the *fully mixed Nash equilibrium* where each agent selects each resource with non-zero probability. While offering in addition an advantage with respect to convenience in handling, the fully mixed Nash equilibrium is suitable for our economic framework under the very natural assumption that each resource offers usage to all agents without imposing any access restrictions.

We define *Optimum* as the *least possible* maximum (over resources) Resource Cost; for a Nash equilibrium, we define *Social Cost* as the expectation, over random choices of the agents,

*One could argue that our pricing scheme is *unfair* in the sense that players with smaller demand can be forced to support those players with larger demand that share the same resource. However, we chose to coin our model as fair on account of the fact that it treats all players sharing the same resource equally, while it does not overcharge players beyond the actual cost of the resource they choose.

of maximum (over resources) Resource Cost. We observe that Optimum is achieved when all agents choose the same resource (Proposition 2.1).

The *Price of Anarchy* [11, 19] is the ratio of Social Cost over Optimum in the *worst-case* pair of chosen demands and Nash equilibrium. To evaluate the Nash Equilibria of the FairPricingGame, we use both the Price of Anarchy and the *Diffuse Price of Anarchy*, an extension of the former, which we introduce to accommodate the (arbitrary but specific) probability distribution on the demands.

1.2 The Diffuse Price of Anarchy

The main argument for using worst-case demands in the definition of the Price of Anarchy [11] is that the distribution of the demands is not known. However, the Price of Anarchy takes this argument way too far. It assumes that nothing is known about the distribution, so that any distribution on the demands is possible. The worst-case distribution prevailing in the definition of the Price of Anarchy is the one in which the worst-case demand occurs with probability one. We extend this definition to remove this assumption while avoiding to assume full knowledge about the distribution of demands. Roughly speaking, the *Diffuse Price of Anarchy* is the *worst-case*, over all allowed probability distributions, of the expectation (according to each specific probability distribution) of the ratio of Social Cost over Optimum in the worst-case Nash equilibrium. It is hoped that the Diffuse Price of Anarchy will alleviate the often overly pessimistic, original notion of the Price of Anarchy [11].

1.3 Summary and Contribution

We prove that the FairPricingGame may not have a pure Nash equilibrium unless all chosen demands are identical (Theorem 3.1). The proof argues that the Resource Costs of all resources are identical in a pure Nash equilibrium. This result is fairly intuitive: when players have different demands, "heavy" players seek to share their costs with "light" players, while "light" players want to avoid "heavy" players.

On the other hand, we show that, for the Fair Pricing model, a fully mixed Nash equilibrium always exists (Theorem 4.1). For the case of two agents, we prove that the fully mixed Nash equilibrium is the unique Nash equilibrium (Theorem 4.2).

We next prove that the Price of Anarchy can be as large as $\Theta(n)$; we establish both lower and upper bounds (Theorems 5.1 and 5.3). A tighter analysis applies to the case of two agents to yield that the Price of Anarchy is then less than $2 - \frac{1}{m}$ (Theorem 5.2).

To mitigate the high $\Theta(n)$ bound on the Price of Anarchy, we seek structure in the probability distributions of demands. The outcome has been the identification of an interesting

class of probability distributions on demands for which the Diffuse Price of Anarchy is upper bounded by a constant in the asymptotic case of large games with a number of players that goes to infinity (Theorem 6.3 and Corollary 6.4). This is a very simple but broad class of so called *bounded, independent probability distributions* – there, roughly speaking, demands are independent and identically distributed, while each demand may not exceed some *universal* for the class constant times its expectation. Under the additional assumption that each demand is symmetrically distributed around its expectation, this universal constant is just 2. To the best of our knowledge, our work is the first to prove something nontrivial about an averaged Price of Anarchy such as the Diffuse Price of Anarchy we have introduced.

Studied in the literature on Economic Theory are several cost sharing models similar to (but different than) our Fair Pricing model. Such cost sharing models have adopted definitions for Resource Cost and Individual Cost similar to the ones we employed in our Fair Pricing model. We conclude this article by considering (within the framework adopted in this work) two prominent examples of such cost sharing models, namely the *Average Cost Pricing* [2] and *Serial Cost Sharing* [16, 17] models. In contrast to our Fair Pricing model, we establish the existence of Pure Nash equilibria for both these models (Theorems 7.1 and 7.2, respectively).

1.4 Related Work

Our Fair Pricing model is partly inspired by the KP model [11] for selfish routing; there is a vast volume of research about selfish routing with unsplittable demands (see, e.g., [3, 4, 5, 10, 11, 13]). The KP model overcharges the total demand on a resource to each and every agent choosing the resource. To the best of our knowledge, our Fair Pricing Model is the first work to explicitly formulate and evaluate, using the Price of Anarchy, a theoretical model of cost sharing with resources and selfish agents that charges the fair share of the total demand to each agent choosing the resource; of course, proportional cost sharing models such as ours have been considered before in Computer Science (e.g., in the microprocessor scheduling/sharing model).

Most cost sharing models studied in the literature on Economic Theory have considered an economic system with a single resource and multiple agents, each choosing its individual demand as its strategy; these models have addressed the problem of identifying the most suitable cost function for the single resource according to several axiomatic criteria, such as monotonicity and envy-freeness. The essential differences between our Fair Pricing model and those models are that we consider multiple resources (albeit identical), and we model the strategy of an agent as some resource. As a result, associated with the two approaches are different notions of equilibria which exhibit different properties.

The Price of Anarchy was originally proposed by Koutsoupias and Papadimitriou [11], and further advocated by Papadimitriou [19], as a measure of performance degradation in systems

with resources shared by selfish agents. The Diffuse Price of Anarchy is motivated by the *Diffuse Adversary* studied by Koutsoupias and Papadimitriou [12] as an alternative to *worst-case* adversaries usually considered for competitive analysis in online computing. So, Diffuse Competitive Ratio is to Competitive Ratio in Online Computing what Diffuse Price of Anarchy is to Price of Anarchy in Selfish Computing.

Making assumptions on the probability distribution of demands (and studying the resulting Diffuse Price of Anarchy) has been our approach towards mitigating the high Price of Anarchy. Coordination mechanisms [3] are another means for reducing the Price of Anarchy.

The fully mixed Nash equilibrium was originally proposed by Mavronicolas and Spirakis [15]; its various existence and uniqueness properties were subsequently studied very extensively; see, e.g., [4, 5, 6, 7, 13, 14].

Anshelevich *et al.* [1] have introduced and studied a *network design* game (and its associated Nash equilibria), where a collection of self-interested (but *unweighted*) agents wish to form a network connecting certain endpoints. Each agent is charged for each edge used for its connection; the cost of each such edge is equally divided among all agents sharing the edge. This *Fair Division* scheme is identical to our Fair Pricing model.

Hayrapetyan *et al.* [8] presented and analyzed (using the Price of Anarchy) a pricing game to capture the interaction between service providers and users over the Internet. Their pricing model addresses the competition of network managers for users via *prices* and the quality of service provided; it addresses neither fairness nor distributions on demands, which are the two main ingredients of our model.

1.5 Organization

Section 2 introduces the Fair Pricing model and summarizes some preliminary facts. Pure Nash equilibria and fully mixed Nash equilibria are treated in Section 3 and 4, respectively. Sections 5 and 6 present our results for the Price of Anarchy and the Diffuse Price of Anarchy, respectively. Section 7 considers some additional cost sharing models. We conclude, in Section 8, with a discussion of our results and some suggestions for further research.

2 The Fair Pricing Model

Our Fair Pricing model was originally motivated by the standard KP model for selfish routing [11]; it departs from it by encompassing some stochastic assumptions on user demands, and notions of pricing and fairness as well.

2.1 Notation

For an event E in a sample space, denote $\Pr\{E\}$ the probability of event E occurring. For a random variable X that follows the probability distribution D , denote $\mathcal{E}_D(X)$ the *expectation* of X (according to the probability distribution D). For any integer $m \geq 2$, denote $[m] = \{1, \dots, m\}$.

2.2 Agents and Resources

We consider a collection $\mathcal{M} = \{1, 2, \dots, m\}$ of identical *resources*, and a collection $\mathcal{N} = \{1, 2, \dots, n\}$ of *agents*. Associated with an agent $i \in \mathcal{N}$ is a *demand* $w_i \in \mathbb{R}_+$. We assume that demands are chosen according to some (known) joint probability distribution D , which comes from some (known) class Δ of possible distributions. We consider D to be the steady state distribution of some ergodic stochastic process that generates demands. We will often fix a particular outcome of the experiment of choosing demands (according to D), which is a $n \times 1$ *demand vector* \mathbf{w} . Denote $W = \sum_{i \in \mathcal{N}} w_i$ and $\widehat{W} = \frac{W}{n}$. Note that $\frac{w_i}{\widehat{W}} \leq n$ for all agents $i \in \mathcal{N}$. For each integer $k \in [n]$, denote $W_k = \sum_{i \in [k]} w_i$; take $W_0 = 0$. (Clearly, $W_n = W$.) For each agent $i \in \mathcal{N}$, denote $W_{-i} = \sum_{k \in [n] \setminus \{i\}} w_k$. We will be assuming, without loss of generality, that $w_1 \geq w_2 \geq \dots \geq w_n$. We usually use subscripts for agents and superscripts for resources.

2.3 Strategies and Assignments

A *pure strategy* for agent $i \in \mathcal{N}$ is some specific resource. A *pure assignment* $\mathbf{L} \in \mathcal{M}^n$ is a collection of pure strategies, one per agent.

A *mixed strategy* for agent i is a probability distribution on the set of pure strategies. A *mixed assignment* \mathbf{P} is a collection of mixed strategies, one per agent. We emphasize that all mixed strategies in a mixed assignment are stochastically independent. A mixed assignment is represented by an $n \times m$ *probability matrix* \mathbf{P} of mn probabilities p_i^j , $i \in \mathcal{N}$ and $j \in \mathcal{M}$, where p_i^j is the probability that agent i selects resource j . Clearly, for each agent $i \in \mathcal{N}$, $\sum_{j \in \mathcal{M}} p_i^j = 1$.

For each agent $i \in \mathcal{N}$, the *support* of agent i in the mixed assignment \mathbf{P} is the set of resources $\mathcal{S}_i = \{j \in \mathcal{M} \mid p_i^j > 0\}$; thus, the support of agent i is the set of resources which i chooses with non-zero probability.

The mixed assignment \mathbf{P} is *fully mixed* [15] if for all agents $i \in \mathcal{N}$ and resources $j \in \mathcal{M}$, $p_i^j > 0$; thus, each agent selects each resource with non-zero probability. The *standard fully mixed* assignment, denote \mathbf{F} , is the assignment in which $f_i^j = \frac{1}{m}$ for all agents $i \in \mathcal{N}$ and resources $j \in [m]$.

2.4 Resource Demand, Resource Congestion and Resource Cost

Fix a pure assignment $\mathbf{L} = \langle l_1, l_2, \dots, l_n \rangle$ and a resource $j \in \mathcal{M}$. Define the *Resource Demand* on resource j , denoted $W^j(\mathbf{L})$, as the total demand on resource j ; that is,

$$W^j(\mathbf{L}) = \sum_{k \in \mathcal{N}: l_k = j} w_k .$$

Define the *Resource Congestion* on resource j , denoted $n^j(\mathbf{L})$, as the total number of agents on resource j ; that is,

$$n^j(\mathbf{L}) = \sum_{k \in \mathcal{N}: l_k = j} 1 = |\{k \in \mathcal{N} \mid l_k = j\}| .$$

The *Resource Cost* on resource j , denoted $\text{RC}^j(\mathbf{L})$, is the ratio $\frac{W^j(\mathbf{L})}{n^j(\mathbf{L})}$ if $n^j(\mathbf{L}) > 0$, and 0 otherwise.

2.5 Individual Cost and Resource Profit

The *Individual Cost* for agent $i \in \mathcal{N}$, denoted $\text{IC}_i(\mathbf{L})$, is defined to be $\text{RC}^{l_i}(\mathbf{L})$; so,

$$\text{IC}_i(\mathbf{L}) = \frac{W^{l_i}(\mathbf{L})}{n^{l_i}(\mathbf{L})} ,$$

and the Individual Cost of agent i is the Resource Cost of the resource it chooses. Although Individual Cost is identified with Resource Cost in the specific case of the Fair Pricing model considered here, this may not be the case in general. So, we chose to introduce both in order to offer convenience to future cost sharing models that will explicitly distinguish between them.

The *Resource Profit* of resource $j \in \mathcal{M}$, denoted $\text{RP}^j(\mathbf{L})$, is defined as

$$\text{RP}^j(\mathbf{L}) = \sum_{k \in \mathcal{N}: l_k = j} \text{IC}_k(\mathbf{L}) - W^j(\mathbf{L}) .$$

intuitively, the profit of a resource is the total cost it charges to agents choosing it minus the total demand it serves. This definition is very general since it applies to all possible specifications of Individual Cost. Clearly, for the specific way we defined Individual Cost in this work, all Resource Profits are zero. However, we still chose to introduce Resource Profit as an important metric for general theoretical models of cost sharing, even though it happens to be zero in the specific case of the Fair Pricing model considered here. We believe that Resource Profit merits explicit investigation in other cost sharing models as well.

We often drop the arguments of the various costs when these are clear from context.

2.6 Expectations in Mixed Assignments

For a mixed assignment \mathbf{P} , all resource demand, resource congestion and Resource Cost become random variables induced by the probability distribution \mathbf{P} . We define the *Expected Resource Demand*, the *Expected Congestion Demand* and the *Expected Resource Cost* as the expectations of Resource Demand, Resource Congestion and Resource Cost, respectively, according to \mathbf{P} .

The *Expected Individual Cost* $\text{IC}_i(\mathbf{P})$ of agent $i \in \mathcal{N}$ is the expectation of its Individual Cost $\text{IC}_i(\mathbf{L})$, where the pure assignment \mathbf{L} is chosen according to the mixed assignment \mathbf{P} ; so,

$$\text{IC}_i(\mathbf{P}) = \sum_{\mathbf{P} \in \mathcal{M}^n} \left(\prod_{i \in \mathcal{N}} p_i^{l_i} \right) \cdot \text{IC}_i(\mathbf{L}).$$

The *Conditional Expected Individual Cost* $\text{IC}_i^j(\mathbf{P})$ of agent $i \in \mathcal{N}$ on resource $j \in \mathcal{M}$ is the conditional expectation of the Individual Cost of agent i had it been assigned to resource j . Clearly,

$$\text{IC}_i(\mathbf{P}) = \sum_{j \in \mathcal{M}} p_i^j \text{IC}_i^j(\mathbf{P}).$$

2.7 Nash Equilibria

The definition of Expected Individual Cost completes the definition of a strategic game that models fair pricing of resource usage, which we call `FairPricingGame`. We are interested in the induced (both pure and mixed) *Nash equilibria* [18] of `FairPricingGame`.

The pure assignment \mathbf{L} is a *pure Nash equilibrium* if for each agent $i \in \mathcal{N}$, the Individual Cost $\text{IC}_i(\mathbf{L})$ is minimized (given the pure strategies of the other agents); thus, no agent can unilaterally improve its own Individual Cost by switching to another pure strategy.

The mixed assignment \mathbf{P} is a *mixed Nash equilibrium* if for each agent $i \in \mathcal{N}$, the Expected Individual Cost $\text{IC}_i(\mathbf{P})$ is minimized (given the mixed strategies of the other agents); thus, no agent can unilaterally improve its own Expected Individual Cost by switching to another mixed strategy.

The particular definition of Expected Individual Cost implies that for each agent $i \in \mathcal{N}$, for each resource $j \in \mathcal{M}$ such that $p_i^j > 0$, all Conditional Expected Individual Costs IC_i^j are the same and no more than any Conditional Expected Individual Cost IC_i^l for a resource $l \in \mathcal{M}$ with $p_i^l = 0$. This implies that the Expected Individual Cost $\text{IC}_i(\mathbf{P})$ of an agent $i \in \mathcal{N}$ is equal to its Conditional Expected Individual Cost $\text{IC}_i^j(\mathbf{P})$ for any resource $j \in \mathcal{M}$ such that $p_i^j > 0$, and no more than its Conditional Expected Individual Cost $\text{IC}_i^l(\mathbf{P})$ for any resource $l \in \mathcal{M}$ such that $p_i^l = 0$.

A *fully mixed Nash equilibrium* [15] is a fully mixed assignment that is a Nash equilibrium. The *standard fully mixed Nash equilibrium* corresponds (when it exists) to the standard fully mixed assignment.

2.8 Social Cost and Optimum

We proceed to define Social Cost, Optimum and the Price of Anarchy for the specific FairPricingGame we consider. Associated with a mixed Nash equilibrium \mathbf{P} is the *Social Cost* $\text{SC}(\mathbf{w}, \mathbf{P})$, which is the expectation, over all random choices of the agents, of the maximum Resource Cost; thus, $\text{SC}(\mathbf{w}, \mathbf{P}) = \mathcal{E}_{\mathbf{P}} (\max_{j \in \mathcal{M}} \text{RC}^j)$. By definition of Resource Cost, we may explicitly write

$$\text{SC}(\mathbf{w}, \mathbf{P}) = \sum_{\langle l_1, l_2, \dots, l_n \rangle \in \mathcal{M}^n} \prod_{i=1}^n p_i^{l_i} \max_{j \in \mathcal{M} \mid |\{k \in \mathcal{N} \mid l_k = j\}| > 0} \left\{ \frac{\sum_{k \in \mathcal{N} \mid l_k = j} w_k}{|\{k \in \mathcal{N} \mid l_k = j\}|} \right\}.$$

On the other hand, the *Optimum* associated with a demand vector \mathbf{w} , denoted $\text{OPT}(\mathbf{w})$, is the *least possible*, over all pure assignments, maximum Resource Cost; thus,

$$\begin{aligned} \text{OPT}(\mathbf{w}) &= \min_{\mathbf{L} \in \mathcal{M}^n} \max_{j \in \mathcal{M}} \text{RC}^j(\mathbf{w}, \mathbf{L}) \\ &= \min_{\mathbf{L} \in \mathcal{M}^n} \text{SC}(\mathbf{w}, \mathbf{L}), \end{aligned}$$

and explicitly,

$$\text{OPT}(\mathbf{w}) = \min_{\langle l_1, l_2, \dots, l_n \rangle \in \mathcal{M}^n} \max_{j \in \mathcal{M} \mid |\{k \in \mathcal{N} \mid l_k = j\}| > 0} \left\{ \frac{\sum_{k \in \mathcal{N} \mid l_k = j} w_k}{|\{k \in \mathcal{N} \mid l_k = j\}|} \right\}.$$

Clearly, there is always a pure assignment achieving the Optimum. We prove:

Proposition 2.1 *Consider the Fair Pricing model. Then, for any demand vector \mathbf{w} ,*

$$\text{OPT}(\mathbf{w}) = \frac{W}{n}.$$

Proof: Fix any demand vector \mathbf{w} . Clearly, the pure assignment \mathbf{L} where all agents are assigned to the same resource achieves Social Cost $\frac{W}{n}$. Since \mathbf{L} is no better than the optimal assignment, it follows that $\text{OPT}(\mathbf{w}) \leq \frac{W}{n}$. So, it only remains to prove that $\text{OPT}(\mathbf{w}) \geq \frac{W}{n}$.

Consider any arbitrary assignment \mathbf{L} . Consider a resource $l \in [m]$ such that $\text{SC}(\mathbf{w}, \mathbf{L}) = \frac{W^l(\mathbf{L})}{n^l(\mathbf{L})}$; clearly, $n^l(\mathbf{L}) > 0$ and $W^l(\mathbf{L}) > 0$. By definition of Social Cost, it follows that for any resource $j \in [m]$ such that $n^j(\mathbf{L}) > 0$ (and $W^j(\mathbf{L}) > 0$),

$$\frac{W^l(\mathbf{L})}{n^l(\mathbf{L})} \geq \frac{W^j(\mathbf{L})}{n^j(\mathbf{L})}$$

or

$$\frac{n^j(\mathbf{L})}{n^l(\mathbf{L})} \geq \frac{W^j(\mathbf{L})}{W^l(\mathbf{L})}.$$

Summing up over all such resources $j \in [m]$ yields that

$$\sum_{j \in [m] : n^j(\mathbf{L}) > 0} \frac{n^j(\mathbf{L})}{n^l(\mathbf{L})} \geq \sum_{j \in [m] : n^j(\mathbf{L}) > 0} \frac{W^j(\mathbf{L})}{W^l(\mathbf{L})}$$

or

$$\frac{\sum_{j \in [m] : n^j(\mathbf{L}) > 0} n^j(\mathbf{L})}{n^l(\mathbf{L})} \geq \frac{\sum_{j \in [m] : n^j(\mathbf{L}) > 0} W^j(\mathbf{L})}{W^l(\mathbf{L})}.$$

Hence, it follows that $\frac{W^l(\mathbf{L})}{n^l(\mathbf{L})} \geq \frac{W}{n}$. By choice of resource l , this implies that $\text{SC}(\mathbf{w}, \mathbf{L}) \geq \frac{W}{n}$. Since \mathbf{L} was chosen arbitrarily, it follows that $\min_{\mathbf{L} \in \mathcal{M}^m} \text{SC}(\mathbf{w}, \mathbf{L}) \geq \frac{W}{n}$, or $\text{OPT}(\mathbf{w}) \geq \frac{W}{n}$, as needed to complete the proof. \blacksquare

2.9 Price of Anarchy and Diffuse Price of Anarchy

The *Price of Anarchy* (also referred to as *Coordination Ratio* [11]), denoted \mathcal{PA} , is the maximum value, over all demand vectors \mathbf{w} and (mixed) Nash equilibria \mathbf{P} , of the ratio $\frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})}$. Proposition 2.1 immediately implies that

$$\mathcal{PA} = \max_{\mathbf{w}, \mathbf{P}} \left(\frac{n}{W} \cdot \text{SC}(\mathbf{w}, \mathbf{P}) \right).$$

The *Diffuse Price of Anarchy* for the class Δ is given by

$$\mathcal{DPA}_\Delta = \max_{D \in \Delta} \left(\mathcal{E}_D \left(\max_{\mathbf{P}} \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \right) \right).$$

Note that each demand vector \mathbf{w} is a random variable chosen according to the probability distribution D . So, each fixed but arbitrary demand vector \mathbf{w} (chosen according to \mathbf{P}) induces a candidate value $\max_{\mathbf{P}} \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})}$ for the Price of Anarchy (corresponding to the worst-case Nash equilibrium associated with \mathbf{w}), which is a function of the particular demand vector. In turn, each fixed but arbitrary probability distribution D induces an expectation on this value, which is a function of the particular probability distribution D . Finally, the maximum of these expectations, over all possible probability distributions D in the class of probability distributions Δ , is the Diffuse Price of Anarchy.

Proposition 2.1 immediately implies that

$$\mathcal{DPA}_\Delta = \max_{D \in \Delta} \left(\mathcal{E}_D \left(\frac{n}{W} \left(\max_{\mathbf{P}} \text{SC}(\mathbf{w}, \mathbf{P}) \right) \right) \right).$$

3 Pure Nash Equilibria

We prove:

Theorem 3.1 (Pure Nash Equilibria for Fair Pricing) *Consider the Fair Pricing model. Then, there is a pure Nash equilibrium if and only if all demands are identical.*

Proof: Assume first that all demands are equal to w ; that is, all demands are identical. Then, in any pure assignment, the Resource Cost on each resource j such that $n^j > 0$ is equal to w , which implies that all Individual Costs are also equal to w . Hence, every pure assignment is a Nash equilibrium.

Assume now that demands are not necessarily all identical, while there is a pure Nash equilibrium \mathbf{L} . For each resource $l \in \mathcal{M}$, denote $w_1^l, \dots, w_{n^l}^l$ the demands assigned to resource l . So, $\sum_{1 \leq k \leq n^l} w_k^l = W^l$. Fix now a resource $j \in \mathcal{M}$ with $n^j > 0$. Since \mathbf{L} is a Nash equilibrium, for each agent $k \in \{1, 2, \dots, n^j\}$ assigned to resource j , and for each resource $l \in \mathcal{M}$, $l \neq j$, it holds that $\text{IC}_k^j \leq \text{IC}_k^l$ or

$$\frac{W^j}{n^j} \leq \frac{W^l + w_k^j}{n^l + 1};$$

call this the *Nash inequality* for agent k . Summing up over all such agents k yields that

$$W^j \leq \frac{n^j W^l}{n^l + 1} + \frac{W^j}{n^l + 1}.$$

Rearranging terms yields that

$$n^l W^j \leq n^j W^l.$$

This implies that for any resource $l \in \mathcal{M}$ with $n^l > 0$,

$$\frac{W^j}{n^j} \leq \frac{W^l}{n^l}.$$

Since both $n^j, n^l > 0$, a symmetric reasoning applies to yield that for each resource $l \in \mathcal{M}$ with $n^l > 0$,

$$\frac{W^j}{n^j} = \frac{W^l}{n^l}.$$

Note now that for each agent $k \in \{1, 2, \dots, n^j\}$ and for each resource $l \in \mathcal{M}$, $l \neq j$, the Nash inequality for agent k yields that

$$\frac{w_k^j}{n^l + 1} \geq \frac{W^j}{n^j} - \frac{W^l}{n^l + 1}.$$

We consider the implications of this inequality in two possible cases.

- Assume first that $n^l = 0$, so that $W^l = 0$ as well. It follows that $w_k^j \geq \frac{W^j}{n^j}$.
- Assume now that $n^l > 0$. In this case, recall that $\frac{W^j}{n^j} = \frac{W^l}{n^l}$. So, the Nash inequality implies that

$$\begin{aligned} \frac{w_k^j}{n^l + 1} &\geq \frac{W^l}{n^l} - \frac{W^l}{n^l + 1} \\ &= \frac{W^l}{n^l(n^l + 1)}, \end{aligned}$$

which implies that

$$\begin{aligned} w_k^j &\geq \frac{W^l}{n^l} \\ &= \frac{W^j}{n^j}. \end{aligned}$$

So, in all cases, $w_k^j \geq \frac{W^j}{n^j}$ for all agents $k \in \{1, 2, \dots, n^j\}$. This implies that $w_1^j = w_2^j = \dots = w_{n^j}^j = \frac{W^j}{n^j}$; thus, all demands on resource j are identical. Since, however, $\frac{W^j}{n^j} = \frac{W^l}{n^l}$ for any resource $l \in \mathcal{M}$ with $n^l > 0$, and demands are only assigned to such resources, it follows that all demands are identical, as needed. ■

4 Fully Mixed Nash Equilibria

We prove:

Theorem 4.1 (Existence of a Fully Mixed Nash Equilibrium) *Consider the Fair Pricing model. Then, there is a fully mixed Nash equilibrium.*

Proof: Fix any demand vector \mathbf{w} . Consider the standard fully mixed assignment \mathbf{F} . Fix any arbitrary pair of an agent $i \in \mathcal{N}$ and a resource $j \in \mathcal{M}$. We calculate the Conditional Expected Individual Cost of agent i on resource j . There are two cases.

- Assume first that no agent other than i chooses resource j . This occurs with probability $\left(1 - \frac{1}{m}\right)^{n-1}$, and it contributes $w_i \left(1 - \frac{1}{m}\right)^{n-1}$ to IC_i^j .
- Fix now any integer k , where $2 \leq k \leq n$, and assume that $k-1$ agents other than i choose the resource j . This occurs with probability $\left(\frac{1}{m}\right)^{k-1} \left(1 - \frac{1}{m}\right)^{n-k}$. There are $\binom{n-1}{k-1}$ pure assignments where exactly $k-1$ agents (other than agent i) choose resource j , and each

such agent $t \neq i$ selects resource j in exactly $\binom{n-2}{k-2}$ of these assignments. Thus, the total contribution of all such assignments to IC_i^j is

$$\frac{1}{k} \cdot \left(\frac{1}{m}\right)^{k-1} \left(1 - \frac{1}{m}\right)^{n-k} \cdot \left(\binom{n-1}{k-1} w_i + \binom{n-2}{k-2} W_{-i} \right).$$

Hence, it follows that

$$\text{IC}_i^j(\mathbf{F}) = w_i \left(1 - \frac{1}{m}\right)^{n-1} + \sum_{k=2}^n \frac{1}{k} \left(\frac{1}{m}\right)^{k-1} \left(1 - \frac{1}{m}\right)^{n-k} \left(\binom{n-1}{k-1} w_i + \binom{n-2}{k-2} W_{-i} \right).$$

Since $\text{IC}_i^j(\mathbf{F})$ is independent of j , the definition of Nash equilibrium implies that \mathbf{F} is a Nash equilibrium, as needed. \blacksquare

We note that the proof of Theorem 4.1 establishes that the standard fully mixed assignment \mathbf{F} is a Nash equilibrium. Call it the *standard fully mixed Nash equilibrium*. We next present a combinatorial proof of a uniqueness property for the standard fully mixed Nash equilibrium in the case of two agents.

Theorem 4.2 (A Unique Nash Equilibrium for Two Agents) *Consider the Fair Pricing model. Then, the standard fully mixed Nash equilibrium is the unique Nash equilibrium in the case of two agents with nonidentical demands.*

Proof: Fix any demand vector \mathbf{w} . Consider any arbitrary Nash equilibrium \mathbf{P} . We will prove that (necessarily) $\mathbf{P} = \mathbf{F}$, the standard fully mixed Nash equilibrium. Recall that the Expected Individual Cost $\text{IC}_i(\mathbf{P})$ of an agent $i \in \mathcal{N}$ is equal to its Conditional Expected Individual Cost $\text{IC}_i^j(\mathbf{P})$ for any resource $j \in \mathcal{M}$ such that $p_i^j > 0$, and no more than its Conditional Expected Individual Cost $\text{IC}_i^l(\mathbf{P})$ for any resource $l \in \mathcal{M}$ such that $p_i^l = 0$.

We first prove a simple fact, namely that the supports \mathcal{S}_1 and \mathcal{S}_2 of agents 1 and 2, respectively, intersect.

By way of contradiction, assume otherwise; that is, assume that $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$. Clearly, the assumption implies that $\text{IC}_1(\mathbf{P}) = w_1$. Without loss of generality, take that $w_1 > w_2$. Consider any resource $\ell \in \mathcal{S}_2$. Clearly,

$$\begin{aligned} & \text{IC}_1(\mathbf{P}) \\ &= w_1 \\ &> w_1 + \frac{w_2 - w_1}{2} p_2^\ell \quad (\text{since } w_1 > w_2) \\ &= w_1(1 - p_2^\ell) + \frac{w_1 + w_2}{2} p_2^\ell \\ &= \text{IC}_1^\ell(\mathbf{P}), \end{aligned}$$

a contradiction to the fact that \mathbf{P} is a Nash equilibrium.

It follows that $\mathcal{S}_1 \cap \mathcal{S}_2 \neq \emptyset$. So, take any resource $j \in \mathcal{S}_1 \cap \mathcal{S}_2$. Clearly,

$$\begin{aligned} \text{IC}_1(\mathbf{P}) &= \text{IC}_1^j(\mathbf{P}) \\ &= w_1(1 - p_2^j) + \frac{w_1 + w_2}{2} p_2^j \\ &= w_1 + \frac{w_2 - w_1}{2} p_2^j; \end{aligned}$$

similarly, we obtain that

$$\text{IC}_2(\mathbf{P}) = w_2 + \frac{w_1 - w_2}{2} p_1^j.$$

The derived expressions imply that $\text{IC}_1(\mathbf{P}) < w_1$ and $\text{IC}_2(\mathbf{P}) > w_2$. We next argue that $\mathcal{S}_1 = \mathcal{S}_2 = \mathcal{M}$.

- Assume that there exists a resource $k \in \mathcal{S}_1 \setminus \mathcal{S}_2$. Then $\text{IC}_1^k(\mathbf{P}) = w_1 > \text{IC}_1(\mathbf{P})$, a contradiction.
- Assume that there exists a resource $k \in \mathcal{S}_2 \setminus \mathcal{S}_1$. Then $\text{IC}_2^k(\mathbf{P}) = w_2 < \text{IC}_2(\mathbf{P})$, a contradiction.

It follows that $\mathcal{S}_1 = \mathcal{S}_2$. Assume, by way of contradiction, that there is a resource $k \notin \mathcal{S}_1$. Then, $\text{IC}_2^k(\mathbf{P}) = w_2 < \text{IC}_2(\mathbf{P})$, a contradiction. It follows that $\mathcal{S}_1 = \mathcal{S}_2 = \mathcal{M}$.

Fix now any pair of resources $j, k \in \mathcal{S}_1 = \mathcal{M}$. Since \mathbf{P} is a Nash equilibrium, $\text{IC}_1^j(\mathbf{P}) = \text{IC}_1^k(\mathbf{P})$, or equivalently

$$w_1(1 - p_2^j) + \frac{w_1 + w_2}{2} p_2^j = w_1(1 - p_2^k) + \frac{w_1 + w_2}{2} p_2^k,$$

or

$$(w_2 - w_1)p_2^j = (w_2 - w_1)p_2^k,$$

implying that $p_2^j = p_2^k$. Since $\mathcal{S}_2 = \mathcal{M}$, it follows that $p_2^j = \frac{1}{m}$ for each resource $j \in \mathcal{M}$. Similarly, we can prove that $p_1^j = \frac{1}{m}$ for each resource $j \in \mathcal{M}$. So, $\mathbf{P} = \mathbf{F}$, as needed. ■

The assumption of nonidentical demands in Theorem 4.2 is *necessary* since every assignment is a Nash equilibrium when the two demands are identical. Moreover, it does *not* hold in general that the standard fully mixed Nash equilibrium is the unique Nash equilibrium. Consider, for example, the case where $n = 3$ and $m = 2$ with $w_1 = w_2$; it is easy to see that the mixed assignment in which agent 1 (resp., agent 2) is assigned to resource 1 (resp., resource 2), while agent 3 is assigned to each resource with probability $\frac{1}{2}$ is a Nash equilibrium (other than the standard fully mixed). We conjecture, however, that the standard fully mixed Nash equilibrium

is *always* the unique fully mixed Nash equilibrium. Also, it is worth investigating the structure of Nash equilibria for the case where all demands are *distinct*. Is it possible that in this case *every* Nash equilibrium is fully mixed? (If our conjecture about fully mixed Nash equilibria is valid, this is equivalent to asking whether or not, for the case of distinct demands, the standard fully mixed Nash equilibrium is the *unique* Nash equilibrium.)

5 The Price of Anarchy

Lower and upper bounds on Price of Anarchy are presented in Sections 5.1 and 5.2.

5.1 Lower Bound

We prove:

Theorem 5.1 *Consider the Fair Pricing model. Then,*

$$\mathcal{PA} \geq \frac{nw_1}{W} \left(\frac{m-1}{m} \right)^{n-1}.$$

Proof: Fix any demand vector \mathbf{w} . Consider the standard fully mixed Nash equilibrium \mathbf{F} , where each pure assignment occurs with the same probability $\left(\frac{1}{m}\right)^n$.

Note that there are $m(m-1)^{n-1}$ pure assignments in which the agent with maximum demand, say w_1 , is the unique agent assigned to the resource it selects. In these pure assignments, the resource chosen by agent 1 is the resource with maximum Resource Cost. Thus,

$$\begin{aligned} \text{SC}(\mathbf{w}, \mathbf{F}) &\geq \left(\frac{1}{m}\right)^n (m(m-1)^{n-1}w_1) \\ &= \left(\frac{m-1}{m}\right)^{n-1} w_1. \end{aligned}$$

Proposition 2.1 implies that $\text{OPT}(\mathbf{w}) = \frac{W}{n}$. Hence,

$$\mathcal{PA} \geq \frac{nw_1}{W} \left(\frac{m-1}{m} \right)^{n-1},$$

as needed. ■

We note that for large values of m , the lower bound on Price of Anarchy from Theorem 5.1 approaches $\frac{nw_1}{W}$. We will soon prove an upper bound matching this lower bound in the limit.

5.2 Upper Bounds

We first prove an upper bound for the special case of two agents.

Theorem 5.2 *Consider the Fair Pricing model. Assume that $n = 2$. Then,*

$$\mathcal{PA} < 2 - \frac{1}{m}.$$

Proof: Fix any demand vector \mathbf{w} . If $w_1 = w_2 = w$, then any (pure or mixed) assignment has Social Cost w , which is equal to Optimum. In particular, any Nash equilibrium does so, which implies that $\mathcal{PA} = 1$. So take that $w_1 > w_2$, and consider the fully mixed Nash equilibrium \mathbf{F} from Theorem 4.1, which, by Proposition 4.2, is the unique Nash equilibrium in this case.

Note that each pure assignment occurs with the same probability $(\frac{1}{m})^2$. Among all m^2 pure assignments, there are $m(m-1)$ pure assignments for which the maximum Resource Cost is w_1 (occurring when the two demands are put on different resources), and m pure assignments for which the maximum Resource Cost is $\frac{w_1 + w_2}{2}$ (occurring when both demands are put on the same resource). So,

$$\begin{aligned} \text{SC}(\mathbf{w}, \mathbf{F}) &= \left(\frac{1}{m}\right)^2 \left(m(m-1)w_1 + m\frac{w_1 + w_2}{2}\right) \\ &= w_1 \left(1 - \frac{1}{m}\right) + \frac{w_1 + w_2}{2} \frac{1}{m}. \end{aligned}$$

Since

$$\text{OPT}(\mathbf{w}) = \frac{w_1 + w_2}{2},$$

it follows that

$$\begin{aligned} \mathcal{PA} &= \frac{2w_1 \left(1 - \frac{1}{m}\right) + \frac{w_1 + w_2}{2}}{w_1 + w_2} + \frac{1}{m} \\ &< 2 - \frac{2}{m} + \frac{1}{m} \quad (\text{since } m \geq 2) \\ &= 2 - \frac{1}{m}, \end{aligned}$$

as needed. ■

We now proceed to the general case of $n \geq 2$ agents. We prove:

Theorem 5.3 *Consider the Fair Pricing model. Then,*

$$\mathcal{PA} \leq \frac{nw_1}{W}.$$

Proof: Fix any demand vector \mathbf{w} . Note that for any pure assignment and any resource $j \in \mathcal{M}$ such that $n^j > 0$, $\frac{W^j}{n^j} \leq w_1$. So, for any Nash equilibrium \mathbf{P} , $\text{SC}(\mathbf{w}, \mathbf{P}) = \mathcal{E}_{\mathbf{P}} \left(\max_{j \in \mathcal{M}} \frac{W^j}{n^j} \right) \leq w_1$. So the Price of Anarchy is

$$\begin{aligned} \mathcal{PA} &\leq \frac{w_1}{\frac{W}{n}} \\ &= \frac{nw_1}{W}, \end{aligned}$$

as needed. ■

Theorem 5.3 immediately implies:

Corollary 5.4 *Assume that for all agents $k \in \mathcal{N}$, $w_k \leq c \cdot \min_i w_i$, for some constant $c > 1$. Then, $\mathcal{PA} \leq c$.*

6 The Diffuse Price of Anarchy

We prove an upper bound on the Diffuse Price of Anarchy for a special case of the class Δ of probability distributions that generate the demands. We start by defining this special class.

Definition 6.1 (Bounded, Independent Probability Distributions) *The class Δ of bounded, independent probability distributions includes all probability distributions D for which the demands w_i , $i \in \mathcal{N}$, are independent, identically distributed random variables such that:*

- *There is some parameter $\delta_D(n) < \infty$ such that $w_i \in [0, \delta_D(n)]$ for each $i \in \mathcal{N}$.*
- *There is some (universal) constant $\ell_\Delta > 0$ such that $\frac{\delta_D(n)}{\mathcal{E}_D(w_i)} \leq \ell_\Delta$ for each $i \in \mathcal{N}$.*

In our proof, we will use the following general version of *Hoeffding bound* [9].

Proposition 6.1 (Hoeffding Bound [9]) *Let X_1, \dots, X_n be independent random variables with $a_k \leq X_k \leq b_k$, for suitable constants a_k and b_k , for each $k \in [n]$. Denote $X = \sum_{k \in [n]} X_k$. Then, for any $t > 0$,*

$$\Pr \{X - \mathcal{E}(X) \leq -nt\} \leq \exp \left(\frac{-2n^2 t^2}{\sum_{k=1}^n (b_k - a_k)^2} \right).$$

We now prove a simple consequence of Proposition 6.1.

Corollary 6.2 Let X_1, \dots, X_n be independent random variables with $0 \leq X_k \leq \delta(n)$, for some suitable $\delta(n) > 0$, for each $k \in [n]$. Assume that $\mathcal{E}(X_k) = \mu$ for each $k \in [n]$. Denote $X = \sum_{k \in [n]} X_k$. Then, for any constant ε , where $0 < \varepsilon < 1$,

$$\Pr \{X \leq (1 - \varepsilon)n\mu\} \leq \exp\left(\frac{-2n\varepsilon^2\mu^2}{\delta^2(n)}\right).$$

Proof: Note that by linearity of expectation, $\mathcal{E}(X) = \sum_{k \in [n]} \mathcal{E}(X_k) = n\mu$. Fix any arbitrary constant ε with $0 < \varepsilon < 1$, and set $t = \varepsilon\mu$. Proposition 6.1 yields that

$$\Pr \{X - n\mu \leq -n\varepsilon\mu\} \leq \exp\left(\frac{-2n^2\varepsilon^2\mu^2}{\sum_{k \in [n]}(\delta(n) - 0)^2}\right)$$

or, equivalently,

$$\Pr \{X \leq (1 - \varepsilon)n\mu\} \leq \exp\left(\frac{-2n\varepsilon^2\mu^2}{\delta^2(n)}\right),$$

as needed. ■

We are now ready to prove:

Theorem 6.3 Consider the Fair Pricing model and the class Δ of bounded, independent probability distributions. Then, for any constant ε with $0 < \varepsilon < 1$,

$$\mathcal{DPA}_\Delta \leq \frac{\ell_\Delta}{1 - \varepsilon} + n \cdot \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right).$$

Proof: Fix any arbitrary probability distribution $D \in \Delta$. Fix any arbitrary demand vector \mathbf{w} chosen according to D . Denote $\mu_D = \mathcal{E}(w_i)$, the common expectation of all demands w_i , $i \in \mathcal{N}$. Fix an arbitrary Nash equilibrium \mathbf{P} . Then,

$$\begin{aligned} & \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \\ & \leq \frac{n \max_{i \in [n]} w_i}{W} \quad (\text{by Theorem 5.3}) \\ & \leq \frac{n\delta_D(n)}{W} \quad (\text{by definition of } \Delta). \end{aligned}$$

We can also derive that

$$\begin{aligned} & \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \\ & \leq \frac{n \max_{i \in [n]} w_i}{W} \quad (\text{by Theorem 5.3}) \\ & \leq n \quad (\text{since } \max_{i \in [n]} w_i \leq W). \end{aligned}$$

Fix now any arbitrary constant ε with $0 < \varepsilon < 1$. Then,

$$\begin{aligned}
\mathcal{E}_D \left(\max_{\mathbf{P}} \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \right) &\leq \frac{n\delta_D(n)}{(1-\varepsilon)n\mu_D} \cdot \Pr\{W > (1-\varepsilon)n\mu_D\} + n \cdot \Pr\{W \leq (1-\varepsilon)n\mu_D\} \\
&= \frac{\delta_D(n)}{(1-\varepsilon)\mu_D} \cdot \Pr\{W > (1-\varepsilon)n\mu_D\} + n \cdot \Pr\{W \leq (1-\varepsilon)n\mu_D\} \\
&\leq \frac{\ell_\Delta}{1-\varepsilon} \cdot \Pr\{W > (1-\varepsilon)n\mu_D\} + n \cdot \Pr\{W \leq (1-\varepsilon)n\mu_D\} \\
&\leq \frac{\ell_\Delta}{1-\varepsilon} + n \cdot \Pr\{W \leq (1-\varepsilon)n\mu_D\}.
\end{aligned}$$

Now, Corollary 6.2 implies that

$$\begin{aligned}
\Pr\{W \leq (1-\varepsilon)n\mu_D\} &\leq \exp\left(\frac{-2n\varepsilon^2\mu_D^2}{\delta_D^2(n)}\right) \\
&\leq \exp\left(\frac{-2n\varepsilon^2\mu_D^2}{\ell_\Delta^2\mu_D^2}\right) \\
&= \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right).
\end{aligned}$$

Hence, it follows that

$$\mathcal{E}_D \left(\max_{\mathbf{P}} \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \right) \leq \frac{\ell_\Delta}{1-\varepsilon} + n \cdot \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right).$$

Since D was chosen arbitrarily, this implies that

$$\begin{aligned}
\mathcal{DPA}_\Delta &= \max_{D \in \Delta} \mathcal{E}_D \left(\max_{\mathbf{P}} \frac{\text{SC}(\mathbf{w}, \mathbf{P})}{\text{OPT}(\mathbf{w})} \right) \\
&\leq \frac{\ell_\Delta}{1-\varepsilon} + n \cdot \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right),
\end{aligned}$$

as needed. ■

Note that

$$\lim_{n \rightarrow \infty} \left(n \cdot \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right) \right) = 0.$$

Hence, by taking n to tend to infinity and choosing ε to be arbitrarily close to 0, Theorem 6.3 immediately implies an asymptotic upper bound on the Diffuse Price of Anarchy:

Corollary 6.4 *Consider the Fair Pricing model and the class Δ of bounded, independent probability distributions. Then,*

$$\mathcal{DPA}_\Delta \leq \ell_\Delta + o(1).$$

We note that the term $o(1)$ in the statement of Corollary 6.4 represents a sequence that goes to 0 as n goes to infinity.

A significant subclass of the class Δ of bounded, independent probability distributions is the class $\Delta_{sym} \subseteq \Delta$ of *bounded, independent, expectation-symmetric probability distributions*. For each such distribution $D \in \Delta_{sym}$, each demand is distributed symmetrically around its expectation; this happens, for example, when each demand is *uniformly* distributed in the interval $[0, \delta_D(n)]$. In this case, for each demand w_i , $i \in \mathcal{N}$, $\mathcal{E}_D(w_i) = \frac{\delta_D(n)}{2}$. So, in this case, $\ell_{\Delta_{sym}} = 2$, and Theorem 6.3 implies:

Corollary 6.5 *Consider the class Δ_{sym} of bounded, independent, expectation-symmetric probability distributions. Fix an arbitrary constant ε with $0 < \varepsilon < 1$. Then,*

$$\mathcal{DPA}_{\Delta_{sym}} \leq \frac{2}{1-\varepsilon} + n \cdot \exp\left(\frac{-n\varepsilon^2}{4}\right).$$

7 Other Cost Sharing Mechanisms

In this section, we consider (within the framework adopted in this work) two most closely related models from the literature on Economic Theory, namely the *Average Cost Pricing* [2] and the *Serial Cost Sharing* [16, 17] models.

Both of these models use a nondecreasing *Cost Function* C^l for each resource $l \in [m]$, and they take as Resource Cost the value $C^l(\sum_{k \in \mathcal{N} | l_k = l} w_k)$. This value is different than the Resource Cost adopted in the Fair Pricing model, even if C^l is the identity function. The crucial difference is that the latter depends on (in particular, decreases with) the Resource Congestion for resource l , while the former ignores Resource Congestion completely. Moreover, both models seek ways to share the Resource Cost $C^l(\sum_{k \in \mathcal{N} | l_k = l} w_k)$ among the selfish agents choosing the resource l .

Recall that for a *concave* function C over $[0, \infty)$, such that $C(0) = 0$, it holds that for all pairs x, y with $0 < x < y$,

$$\frac{C(x)}{x} \geq \frac{C(y)}{y};$$

for a *strictly concave* function C over $[0, \infty]$, such that $C(0) = 0$, it holds that for all pairs x, y with $0 < x < y$,

$$\frac{C(x)}{x} > \frac{C(y)}{y}.$$

7.1 Average Cost Pricing

In the *Average Cost Pricing* model [2], the Individual Cost of agent $i \in \mathcal{N}$ choosing resource $l \in \mathcal{M}$ in the pure assignment \mathbf{L} is

$$IC_i(\mathbf{L}) = \frac{w_i}{\sum_{k \in \mathcal{N} | l_k = l} w_k} C^l \left(\sum_{k \in \mathcal{N} | l_k = l} w_k \right).$$

We prove:

Theorem 7.1 (Pure Nash Equilibria for Average Cost Pricing) *Consider the Average Cost Pricing model with $C^\ell = C$ for each resource $\ell \in \mathcal{M}$, where C is a strictly concave function over $[0, \infty)$ with $C(0) = 0$. Then, there are exactly m pure Nash equilibria.*

Proof: It suffices to prove that a pure assignment \mathbf{L} is a Nash equilibrium if and only if there is a resource $\ell \in \mathcal{M}$ such that for each user $i \in \mathcal{N}$, $l_i = \ell$; that is, a pure assignment is a Nash equilibrium if and only if it assigns all users to the same resource.

Assume first that the pure assignment \mathbf{L} is such that there is a resource $\ell \in \mathcal{M}$ such that for each user $i \in \mathcal{N}$, $l_i = \ell$.

Fix an agent $i \in \mathcal{N}$ and a resource $j \in \mathcal{M}$, $j \neq \ell$. Then,

$$IC_i^\ell(\mathbf{L}) = \frac{w_i}{W} C(W),$$

while

$$IC_i^j(\mathbf{L}) = \frac{w_i}{w_i} C(w_i),$$

Since C is strictly concave and $w_i < W$,

$$\frac{C(W)}{W} < \frac{C(w_i)}{w_i}.$$

It follows that $IC_i^\ell(\mathbf{L}) < IC_i^j(\mathbf{L})$. Since agent i and resource j were chosen arbitrarily, this implies that \mathbf{L} is a Nash equilibrium.

Assume now that \mathbf{L} is a Nash equilibrium. Towards a contradiction, assume that there are distinct agents $i, k \in \mathcal{N}$ such that $l_i \neq l_k$. We analyze users i and k .

- Since \mathbf{L} is a Nash equilibrium, $IC_i^{l_i}(\mathbf{L}) \leq IC_i^{l_k}(\mathbf{L})$, or

$$\frac{w_i}{W^{l_i}} C(W^{l_i}) \leq \frac{w_i}{W^{l_k} + w_i} C(W^{l_k} + w_i),$$

or

$$\frac{C(W^{l_i})}{W^{l_i}} \leq \frac{C(W^{l_k} + w_i)}{W^{l_k} + w_i}.$$

Since C is a strictly concave function, this implies that $W^{l_i} \geq W^{l_k} + w_i$.

- Since \mathbf{L} is a Nash equilibrium, $\text{IC}_k^{l_k}(\mathbf{L}) \leq \text{IC}_k^{l_i}(\mathbf{L})$, or

$$\frac{w_k}{W^{l_i}} \text{C}(W^{l_k}) \leq \frac{w_k}{W^{l_i} + w_k} \text{C}(W^{l_i} + w_k),$$

or

$$\frac{\text{C}(W^{l_k})}{W^{l_k}} \leq \frac{\text{C}(W^{l_i} + w_k)}{W^{l_i} + w_k}.$$

Since C is a strictly concave function, this implies that $W^{l_k} \geq W^{l_i} + w_k$.

Since $W^{l_i} \geq W^{l_k} + w_i$ and $W^{l_k} \geq W^{l_i} + w_k$, it follows that $0 \geq w_i + w_k$. A contradiction. ■

7.2 Serial Cost Sharing

In the *Serial Cost Sharing* model [16, 17], there are intuitive, systematical formulas for the Individual Costs of the agents choosing a particular resource $l \in \mathcal{M}$, which we will first demonstrate for the special case of three agents with demands $w_1 \leq w_2 \leq w_3$.[†] We refer to a pure assignment \mathbf{L} assigning the three agents to resource l . The Individual Cost of agent 1 with the smallest demand w_1 is

$$\text{IC}_1(\mathbf{L}) = \frac{\text{C}^l(3w_1)}{3}.$$

This is similar to (but different than) the fair share of Resource Cost employed in the Fair Pricing model, while it depends on the Resource Congestion (equal to 3) in a way identical to the one in our model. This is not true for the rest of the agents. The Individual Cost of agent 2 is

$$\begin{aligned} \text{IC}_2(\mathbf{L}) &= \frac{\text{C}^l(w_1 + 2w_2) - \text{IC}_1(\mathbf{L})}{2} \\ &= \frac{\text{C}^l(w_1 + 2w_2)}{2} - \frac{\text{C}^l(3w_1)}{6}. \end{aligned}$$

Finally, agent 3 pays for the rest of Resource Cost, and this is calculated to be

$$\begin{aligned} \text{IC}_3(\mathbf{L}) &= \text{C}^l(w_1 + w_2 + w_3) - \text{IC}_2(\mathbf{L}) - \text{IC}_1(\mathbf{L}) \\ &= \text{C}^l(w_1 + w_2 + w_3) - \frac{\text{C}^l(w_1 + 2w_2)}{2} - \frac{\text{C}^l(3w_1)}{6}. \end{aligned}$$

[†]The assumption $w_1 \leq w_2 \leq w_3$ violates the more general assumption $w_1 \geq \dots \geq w_n$ on the ordering of demands from Section 2.2. We chose to adopt it in order to be consistent with the literature on Serial Cost Sharing [16, 17].

Now for the general case where n^l agents with weights $w_1^l \leq \dots \leq w_{n^l}^l$ choose resource $l \in \mathcal{M}$ according to the pure assignment \mathbf{L} , the Individual Cost of agent $i \in [n^l]$ is

$$\begin{aligned} & \text{IC}_i(\mathbf{L}) \\ &= \sum_{k=1}^i \left(\frac{\mathcal{C}^l(w_1^l + \dots + w_{k-1}^l + (n^l + 1 - k)w_k^l)}{n^l + 1 - k} - \frac{\mathcal{C}^l(w_1^l + \dots + w_{k-2}^l + (n^l + 2 - k)w_{k-1}^l)}{n^l + 1 - k} \right). \end{aligned}$$

Note that all Individual Costs depend on Resource Congestion in the Serial Cost Sharing model, but in a way much more involved (but probably noteworthy) than the one in our Fair Pricing model. The Serial Cost Sharing model reflects the natural principle that the Individual Cost of an agent does not depend on other agents' demands that are larger than its own; obviously, this principle is violated in the Fair Pricing model. We prove:

Theorem 7.2 (Pure Nash Equilibria for Serial Cost Sharing) *Consider the Serial Cost Sharing model, with $\mathcal{C}^\ell = \mathcal{C}$ for each resource $\ell \in \mathcal{M}$, where \mathcal{C} is an increasing, concave function over $[0, \infty)$ with $\mathcal{C}(0) = 0$. Then, there is a pure Nash equilibrium.*

Proof: Assume, without loss of generality, that $w_1 \leq w_2 \leq \dots \leq w_n$, and consider the pure assignment \mathbf{L} assigning all agents to the same resource. We will prove that \mathbf{L} is a Nash equilibrium.

Specifically, we prove by induction on $i \in [n]$ that $\text{IC}_i(\mathbf{L}) \leq \mathcal{C}(w_i)$; this will imply that \mathbf{L} is a Nash equilibrium.

For the basis case where $i = 1$, note that

$$\begin{aligned} & \text{IC}_1(\mathbf{L}) \\ &= \frac{\mathcal{C}(nw_1)}{n} \quad (\text{by definition of Serial Cost Sharing}) \\ &\leq \mathcal{C}(w_1) \quad (\text{since } \mathcal{C} \text{ is concave}), \end{aligned}$$

Assume inductively that $\text{IC}_t(\mathbf{L}) \leq \mathcal{C}(w_t)$ for all agents t with $1 \leq t < n$.

For the induction step, consider agent $t + 1$. Then,

$$\begin{aligned} \text{IC}_{t+1}(\mathbf{L}) &= \sum_{k=1}^{t+1} \frac{\mathcal{C}(W_{k-1} + (n+1-k)w_k) - \mathcal{C}(W_{k-2} + (n+2-k)w_{k-1})}{n+1-k} \\ &= \sum_{k=1}^t \frac{\mathcal{C}(W_{k-1} + (n+1-k)w_k) - \mathcal{C}(W_{k-2} + (n+2-k)w_{k-1})}{n+1-k} \\ &\quad + \frac{\mathcal{C}(W_t + (n-t)w_{t+1}) - \mathcal{C}(W_{t-1} + (n+1-t)w_t)}{n-t} \\ &= \text{IC}_t(\mathbf{L}) + \frac{\mathcal{C}(W_t + (n-t)w_{t+1}) - \mathcal{C}(W_t + (n-t)w_t)}{n-t} \\ &\leq \mathcal{C}(w_t) + \frac{\mathcal{C}(W_t + (n-t)w_{t+1}) - \mathcal{C}(W_t + (n-t)w_t)}{n-t}. \end{aligned}$$

If $w_{t+1} = w_t$, then

$$\begin{aligned} \text{IC}_{t+1}(\mathbf{L}) &= \text{IC}_t(\mathbf{L}) \\ &\leq C(w_t) \\ &= C(w_{t+1}). \end{aligned}$$

Now assume that $w_{t+1} > w_t$. Then,

$$\begin{aligned} W_t + (n-t)w_{t+1} &> W_t + (n-t)w_t \\ &> w_t. \end{aligned}$$

Since C is concave, it follows that

$$\frac{C(W_t + (n-t)w_{t+1}) - C(W_t + (n-t)w_t)}{W_t + (n-t)w_{t+1} - W_t - (n-t)w_t} \leq \frac{C(W_t + (n-t)w_{t+1}) - C(w_t)}{W_t + (n-t)w_{t+1} - w_t},$$

or equivalently

$$\frac{C(W_t + (n-t)w_{t+1}) - C(W_t + (n-t)w_t)}{n-t} \leq (w_{t+1} - w_t) \frac{C(W_t + (n-t)w_{t+1}) - C(w_t)}{W_t + (n-t)w_{t+1} - w_t}.$$

Now observe that $W_t + (n-t)w_{t+1} > w_{t+1} > w_t$, implying that

$$\frac{C(W_t + (n-t)w_{t+1}) - C(w_t)}{W_t + (n-t)w_{t+1} - w_t} \leq \frac{C(w_{t+1}) - C(w_t)}{w_{t+1} - w_t}.$$

So

$$\begin{aligned} \text{IC}_{t+1}(\mathbf{L}) &\leq C(w_t) + (w_{t+1} - w_t) \cdot \frac{C(w_{t+1}) - C(w_t)}{w_{t+1} - w_t} \\ &= C(w_t) + C(w_{t+1}) - C(w_t) \\ &= C(w_{t+1}), \end{aligned}$$

and the inductive proof is now complete. ■

Theorem 7.2 implies that concavity suffices for the existence of pure Nash equilibria in the Serial Cost Sharing model; in contrast, Theorem 7.1 relied on the assumption of strong concavity to establish the existence of pure Nash equilibria in the Average Cost Pricing model.

8 Discussion and Directions for Future Research

We proposed here a very intuitive and pragmatic cost mechanism for pricing the competitive usage of resources shared by selfish agents. This mechanism is both distributed and fair. We

presented results for the (pure and mixed) Nash equilibria of the induced strategic game. We also presented bounds for both the Price of Anarchy (considering worst-case demands) and the Diffuse Price of Anarchy (assuming that demands of agents are drawn according to some probability distribution from some wide class).

Our Fair Pricing model provides a concrete first step toward a systematic way of treating such cost mechanisms for pricing the competitive usage of multiple resources. We are currently examining both more general pricing functions and heterogeneous cases of selfish agents. We believe that our proof techniques will be instrumental to obtaining corresponding results for related models and problems. We also believe that our proposed Diffuse Price of Anarchy is of general applicability in congestion games with players' demands drawn according to some known probability distribution. For the Average Cost Pricing and the Serial Cost Sharing model, our work provides less answers than questions it poses.

We conclude this article with some interesting, related open problems:

1. The most obvious open problem is to close the gap between our lower and upper bounds on Price of Anarchy for the Fair Pricing model.
2. Prove or disprove the conjecture that the standard fully mixed Nash equilibrium is the unique fully mixed Nash equilibrium for the Fair Pricing model.
3. A wide avenue of further research concerns the determination of (tight) bounds on the Diffuse Price of Anarchy for the KP model and its variants.
4. Are the assumptions of strong concavity and concavity for the existence of pure Nash equilibria in the Average Cost Pricing and the Serial Cost Sharing models, respectively, necessary?
5. Investigate properties of fully mixed Nash equilibria for the Average Cost Pricing and the Serial Cost Sharing models.
6. Determine (tight) bounds on the Price of Anarchy and the Diffuse Price of Anarchy for the Average Cost Pricing model and the Serial Cost Sharing model.

References

- [1] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler and T. Roughgarden, "The Price of Stability for Network Design with Fair Cost Allocation," *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 295–304, October 2004.
- [2] L. Billera and D. Heath, "Allocation of Shared Costs: A Set of Axioms Yielding a Unique Procedure", *Mathematics of Operations Research*, Vol. 7, pp. 32–39, 1982.
- [3] G. Christodoulou, E. Koutsoupias and A. Nanavati, "Coordination Mechanisms", *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, J. Diaz, J. Karhumäki, A. Lepistö and D. Sannella eds., pp. 345–357, Vol. 3142, Lecture Notes in Computer Science, Springer-Verlag, July 2004.
- [4] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas and P. Spirakis, "The Structure and Complexity of Nash Equilibria for a Selfish Routing Game", *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, P. Widmayer, F. Triguero, R. Morales, M. Hennessy, S. Eidenbenz and R. Conejo eds., pp. 123–134, Vol. 2380, Lecture Notes in Computer Science, Springer-Verlag, July 2002.
- [5] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and M. Rode, "Nash Equilibria in Discrete Routing Games with Convex Latency Functions", *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, J. Diaz, J. Karhumäki, A. Lepistö and D. Sannella eds., pp. 645–657, Vol. 3142, Lecture Notes in Computer Science, Springer-Verlag, July 2004.
- [6] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Extreme Nash Equilibria," *Proceedings of the 8th Italian Conference on Theoretical Computer Science*, C. Blundo and C. Laneve eds., pp. 1–20, Vol. 2841, Lecture Notes in Computer Science, Springer-Verlag, October 2003.
- [7] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Structure and Complexity of Extreme Nash Equilibria", *Theoretical Computer Science*, Special Issue titled "Game Theory Meets Theoretical Computer Science," M. Mavronicolas and S. Abramsky guest eds., Vol. 343, Nos. 1–2, pp. 133–157, October 2005.
- [8] A. Hayrapetyan, É. Tardos and T. Wexler, "A Network Pricing Game for Selfish Traffic", *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing*, pp. 284–291, July 2005.

- [9] W. Hoeffding, "Probability Inequalities for Sums of Bounded Random Variables", *Journal of the American Statistical Association*, Vol. 58, pp. 13–30, 1963.
- [10] E. Koutsoupias, M. Mavronicolas and P. Spirakis, "Approximate Equilibria and Ball Fusion", *Theory of Computing Systems*, Vol. 36, No. 6, pp. 683–693, November 2003.
- [11] E. Koutsoupias and C. H. Papadimitriou, "Worst-Case Equilibria", *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, G. Meinel and S. Tison eds., pp. 404–413, Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, March 1999.
- [12] E. Koutsoupias and C. H. Papadimitriou, "Beyond Competitive Analysis," *SIAM Journal on Computing*, Vol. 30, No. 1, pp. 300–317, April/May 2000.
- [13] T. Lücking, M. Mavronicolas, B. Monien and M. Rode, "A New Model for Selfish Routing," *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science*, V. Diekert and M. Habib eds., pp. 547–558, Vol. 2996, Lecture Notes in Computer Science, Springer-Verlag, March 2004. Also, accepted to *Theoretical Computer Science*.
- [14] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis and I. Vrto, "Which is the Worst-case Nash Equilibrium?," *Proceedings of the 26th International Symposium on Mathematical Foundations of Computer Science*, B. Rovan and P. Vojtás eds., pp. 551–561, Vol. 2747, Lecture Notes in Computer Science, Springer-Verlag, August 2003.
- [15] M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing", *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, July 2001. Also, accepted to *Algorithmica*.
- [16] H. Moulin and S. Shenker, "Serial Cost Sharing", *Econometrica*, Vol. 60, No. 5, pp. 1009–1037, September 1992.
- [17] H. Moulin and S. Shenker, "Average Cost Pricing versus Serial Cost Sharing: An Axiomatic Comparison", *Journal of Economic Theory*, Vol. 64, No. 1, pp. 178–201, October 1994.
- [18] J. F. Nash, "Non-Cooperative Games", *Annals of Mathematics*, Vol. 54, No. 2, pp. 286–295, 1951.
- [19] C. H. Papadimitriou, "Algorithms, Games and the Internet", *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 749–753, July 2001.
- [20] C. Shapiro and H. R. Varian, *Information Rules – A Strategic Guide to the Network Economy*, Harvard Business School Press, 1998.

- [21] M. Spence, “Nonlinear Prices and Welfare”, *Journal of Marketing Research*, Vol. 8, No. 1, pp. 1–18, August 1977.
- [22] H. R. Varian, “Price Discrimination and Social Welfare”, *American Economic Review*, Vol. 75, No. 4, pp. 870–875, September 1985.