

Cost Sharing Mechanisms for Fair Pricing of Resources Usage*

Marios Mavronicolas¹, Panagiota N. Panagopoulou^{2,3} and Paul G. Spirakis^{2,3}

¹ Department of Computer Science, University of Cyprus
Nicosia CY-1678, Cyprus
mavronic@ucy.ac.cy

² Department of Computer Engineering and Informatics, Patras University, Greece

³ Research Academic Computer Technology Institute
N. Kazantzaki Str., Patras University, Rion, Patras GR 26500, Greece
{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 she 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. In contrast, we have been able to prove that pure Nash equilibria do exist for two closely related cost sharing models, namely the *Average Cost Pricing* and the *Serial Cost Sharing* models.
- 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*

* Partially supported by the EU within the 6th Framework Programme under contracts 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS) and 015964 “Algorithmic Principles for Building Efficient Overlay Computers” (AEOLUS), and by the General Secretariat for Research and Technology of the Greek Ministry of Development within the programme PENED 2003. A previous version of this paper appeared at the *1st International Workshop on Internet and Network Economics* (WINE 2005).

and each is at most a (*universal* for the class) constant times its expectation. Then, the Diffuse Price of Anarchy is at most that same constant, which is just 2 when each demand is distributed symmetrically around its expectation.

Keywords. Cost Sharing, Diffuse Price of Anarchy, Fair Pricing, Resources

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 modeling.

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 she 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, [19,20], or the more recent [18] 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* [16] (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, of maximum (over resources) Resource Cost. We

observe that Optimum is achieved when all agents choose the same resource (Proposition 1).

The *Price of Anarchy* [9,17] 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 [9] 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.

1.3 Summary and Contribution

We prove that the FairPricingGame may not have a pure Nash equilibrium unless all chosen demands are identical (Theorem 1). The proof argues that the Resource Costs of all resources are identical in a pure Nash equilibrium. In contrast, we have been able to prove the existence of pure Nash equilibria for two other cost sharing models, namely the *Average Cost Pricing* [1] and the *Serial Cost Sharing* [14,15] models (Theorems 2 and 3).

On the other hand, we show that, for the Fair Pricing model, a fully mixed Nash equilibrium always exists (Theorem 4). For the case of two agents, we prove that the fully mixed Nash equilibrium is the unique Nash equilibrium (Theorem 5).

We next prove that the Price of Anarchy can be as large as $\Theta(n)$; we establish both lower and upper bounds (Theorems 6 and 8). 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 7).

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 (Theorem 9). 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 a *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.

1.4 Related Work

Our Fair Pricing model is partly motivated by the KP model [9] for selfish routing; there is a vast volume of research about selfish routing with unsplittable demands (see, e.g., [2,3,4,8,9,11]). 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).

Studied in the Economics literature are several pricing models similar to (but different than) our Fair Pricing model. These pricing models have mostly 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. Two prominent examples of such pricing models are *Average Cost Pricing* [1] and *Serial Cost Sharing* [14,15]; they will be discussed in more detail in Sect. 2. The essential differences between our Fair Pricing model and those studied in the Economics literature 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 exhibiting different properties.

The Price of Anarchy was originally proposed by Koutsoupias and Papadimitriou, and further advocated by Papadimitriou [17], 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 [10] 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.

The fully mixed Nash equilibrium was originally proposed by Mavronicolas and Spirakis [13]; its various existence and uniqueness properties were subsequently studied very extensively; see, e.g., [3,4,5,11,12]. Coordination mechanisms [2] are another means for reducing the Price of Anarchy.

Hayrapetyan *et al.* [6] 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.

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

2 The Fair Pricing Model

Our Fair Pricing model was originally motivated by the standard KP model for selfish routing [9]; it departs from it by encompassing some stochastic assumptions on user demands, and notions of pricing and fairness as well.

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 $[m] = \{1, \dots, m\}$.

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 $i \in \mathcal{N}$. 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.

Strategies and Assignments. A *pure strategy* for agent $i \in \mathcal{N}$ is some specific resource; a *mixed strategy* for agent i is a probability distribution on the set of pure strategies. A *pure assignment* $\mathbf{L} \in \mathcal{M}^n$ is a collection of pure strategies, one per agent. Similarly, a *mixed assignment* \mathbf{P} is a collection of mixed strategies, one per agent. 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. \mathbf{P} is *fully mixed* [13] 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.

Resource Demand, Resource Congestion, Resource Cost, Individual Cost and Resource Profit. 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. 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 she chooses. Although Individual Cost is identified with Resource Cost in the specific case of the Fair Pricing model considered here, this may not be true 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.

Other Cost Sharing Models. Similar cost sharing models from the Economics literature have considered Resource Costs and Individual Costs similar to the ones we employed in our Fair Pricing model. We will comment (using the terminology adopted in this work) on the two most closely related ones, namely the *Average Cost Pricing* [1] and the *Serial Cost Sharing* [14,15] models. Both of them use a nondecreasing *Cost Function* C^l for each resource $l \in [m]$, and 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 here, 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 .

- In the *Average Cost Pricing* model [1], the Individual Cost of agent $i \in \mathcal{N}$ choosing resource $l \in \mathcal{M}$ in the pure assignment \mathbf{L} is

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

- In the *Serial Cost Sharing* model [14,15], there are intuitive, systematical formulas for the Individual Costs of the agents choosing a 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} . The Individual Cost of agent 1 with the smallest demand is $\text{IC}_1(\mathbf{L}) = \frac{C^l(3w_1)}{3}$. This is similar to (but different than) the fair share of Resource Cost employed in the Fair Pricing model, and 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

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

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

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

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

$$\text{IC}_i = \sum_{k=1}^i \frac{C^l(w_1^l + \dots + w_{k-1}^l + (n^l + 1 - k)w_k^l)}{n^l + 1 - k} - \frac{C^l(w_1^l + \dots + w_{k-2}^l + (n^l + 2 - k)w_{k-1}^l)}{n^l + 1 - k}.$$

All Individual Costs depend on Resource Congestion, 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 principle that the Individual Cost of an agent does not depend on user demands that are larger than its own; this principle is violated in the Fair Pricing model.

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 *Conditional Expected Individual Cost* IC_i^j of agent $i \in \mathcal{N}$ on resource $j \in \mathcal{M}$ is the conditional expectation of the Individual Cost of agent i had she been assigned to resource j . The *Expected Individual Cost* IC_i of agent i is the expectation of her Conditional Expected Individual Cost on a resource; so, $\text{IC}_i = \sum_{j \in \mathcal{M}} p_i^j \text{IC}_i^j$.

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* [16] of **FairPricingGame**. Formally, 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 her own Individual Cost. The mixed assignment \mathbf{P} is a *mixed Nash equilibrium* if for each agent $i \in \mathcal{N}$, the Expected Individual Cost $\text{IC}_i(\mathbf{L})$ is minimized (given the mixed strategies of the other agents); thus, no agent can unilaterally improve her own Expected Individual Cost. 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 with $p_i^l = 0$.

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,

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

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, $\text{OPT}(\mathbf{w}) = \min_{\mathbf{L} \in \mathcal{M}^n} \max_{j \in \mathcal{M}} \text{RC}^j(\mathbf{w}, \mathbf{L})$, and explicitly,

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

Proposition 1. *For any demand vector \mathbf{w} , $\text{OPT}(\mathbf{w}) = \frac{W}{n}$.*

Proof. Fix any demand vector \mathbf{w} . Clearly, the pure assignment σ where all agents are assigned to the same resource achieves Social Cost $\frac{W}{n}$. Since σ 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 α . Let resource $l \in [m]$ be such that $\text{SC}(\mathbf{w}, \alpha) = \frac{W^l(\alpha)}{n^l(\alpha)}$. By definition of Social Cost, it follows that for any resource $j \in [m]$ such that $n^j > 0$,

$$\frac{W^l(\alpha)}{n^l(\alpha)} \geq \frac{W^j(\alpha)}{n^j(\alpha)} \quad \text{or} \quad \frac{n^j(\alpha)}{n^l(\alpha)} \geq \frac{W^j(\alpha)}{W^l(\alpha)}.$$

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

$$\begin{aligned} \sum_{j \in [m]: n^j > 0} \frac{n^j(\alpha)}{n^l(\alpha)} &\geq \sum_{j \in [m]: n^j > 0} \frac{W^j(\alpha)}{W^l(\alpha)} \\ \frac{\sum_{j \in [m]: n^j > 0} n^j(\alpha)}{n^l(\alpha)} &\geq \frac{\sum_{j \in [m]: n^j > 0} W^j(\alpha)}{W^l(\alpha)}. \end{aligned}$$

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

Price of Anarchy and Diffuse Price of Anarchy. The *Price of Anarchy* (also referred to as *Coordination Ratio* [9]), 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 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

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

Each fixed but arbitrary demand vector \mathbf{w} induces a candidate value 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. Now, 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, is the Diffuse Price of Anarchy. Proposition 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

3.1 Fair Pricing Model

For the `FairPricingGame` we defined in the previous section, we prove:

Theorem 1. *There is a pure Nash equilibrium if and only if all demands are identical.*

Proof. Assume first that all demands are equal to w , i.e. all demands are identical. Then, in any pure Nash 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 there is a pure Nash equilibrium \mathbf{L} . For each resource $j \in \mathcal{M}$, denote $w_1^j, \dots, w_{n^j}^j$ the demands assigned to resource j . So, $\sum_{1 \leq k \leq n^j} w_k^j = W^j$. 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 $\frac{W^j}{n^j} \leq \frac{W^l + w_k^j}{n^l + 1}$. 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 pair of resources $j, l \in \mathcal{M}$ with $n^j, n^l > 0$, $\frac{W^j}{n^j} = \frac{W^l}{n^l}$.

Note now that for each agent $k \in \{1, 2, \dots, n^j\}$, $\frac{w_k^j}{n^l + 1} \geq \frac{W^j}{n^j} - \frac{W^l}{n^l + 1}$. We consider the implications of this inequality in two possible cases.

- Assume first that $n^l = 0$ (in which case $W^l = 0$ as well). Then, $w_k^j \geq \frac{W^j}{n^j}$.
- Assume now that $n^l > 0$. In this case, recall that $\frac{W^j}{n^j} = \frac{W^l}{n^l}$. So, the inequality implies that $\frac{w_k^j}{n^l + 1} \geq \frac{W^l}{n^l} - \frac{W^l}{n^l + 1} = \frac{W^l}{n^l(n^l + 1)}$, implying that $w_k^j \geq \frac{W^l}{n^l} = \frac{W^j}{n^j}$.

So, in all cases, $w_k^j \geq \frac{W^j}{n^j}$ for all $k \in \{1, 2, \dots, n^j\}$. This implies that $w_1^j = w_2^j = \dots = w_{n^j}^j = \frac{W^j}{n^j}$. Since, however, $\frac{W^j}{n^j} = \frac{W^l}{n^l}$ for any pair of resources $j, l \in \mathcal{M}$ with $n^j, n^l > 0$, it follows that all demands are identical, as needed. \square

3.2 Average Cost Pricing

Now consider a special case of the Average Cost Pricing model, where $C^\ell(\cdot) = C(\cdot)$ for all $\ell \in \mathcal{M}$ and $C(\cdot)$ is an increasing, *concave* function in $[0, \infty)$ with $C(0) = 0$. Then

Theorem 2. *There are exactly m pure Nash equilibria for the Average Cost Pricing model. Namely, a pure assignment \mathbf{L} is a pure Nash equilibrium if and only if $\ell_i = \ell \forall i \in \mathcal{N}$, for some $\ell \in \mathcal{M}$.*

Proof. Assume that $w_1 \geq w_2 \geq \dots \geq w_n$ and denote $W = \sum_{i=1}^n w_i$. Assigning all agents on the same resource $\ell \in \mathcal{M}$ is a pure Nash equilibrium, since for all $i \in \mathcal{N}$, $j \neq \ell \in \mathcal{M}$ it holds that

$$\frac{C(W)}{W} \leq \frac{C(w_i)}{w_i} \iff \frac{w_i}{W} C(W) \leq \frac{w_i}{w_i} C(w_i) .$$

It remains to prove that there exists no pure Nash equilibrium \mathbf{L} such that $\ell_i \neq \ell_j$ for some $i, j \in \mathcal{N}$. Assume by contradiction that there exists such a pure Nash equilibrium. Then

$$\begin{aligned} \frac{w_i}{W^{\ell_i}} C(W^{\ell_i}) &\leq \frac{w_i}{W^{\ell_j} + w_i} C(W^{\ell_j} + w_i) \\ \frac{C(W^{\ell_i})}{W^{\ell_i}} &\leq \frac{C(W^{\ell_j} + w_i)}{W^{\ell_j} + w_i} \\ W^{\ell_i} &\geq W^{\ell_j} + w_i . \end{aligned}$$

Similarly

$$W^{\ell_j} \geq W^{\ell_i} + w_j ,$$

implying that $w_i + w_j \leq 0$, a contradiction. \square

3.3 Serial Cost Sharing

Similarly, consider a special case of the Serial Cost Sharing model, where $C^\ell(\cdot) = C(\cdot)$ for all $\ell \in \mathcal{M}$ and $C(\cdot)$ is an increasing, *concave* function in $[0, \infty)$ with $C(0) = 0$. Then

Theorem 3. *There exists a pure Nash equilibrium.*

Proof. Let $w_1 \leq w_2 \leq \dots \leq w_n$ and assume that all agents choose the same resource. We will prove by induction on i that $IC_i \leq C(w_i)$, so that this assignment be a pure Nash equilibrium. For $1 \leq t \leq n$ denote $W_t = \sum_{i=1}^t w_i$.

- For $i = 1$, $IC_1 = \frac{C(nw_1)}{n} \leq C(w_1)$, since $C(\cdot)$ is concave.
- Suppose that $IC_t \leq C(w_t)$ for some $1 \leq t < n$.
- Then

$$\begin{aligned} IC_{t+1} &= \sum_{k=1}^{t+1} \frac{C(W_{k-1} + (n+1-k)w_k) - C(W_{k-2} + (n+2-k)w_{k-1})}{n+1-k} \\ &= \sum_{k=1}^t \frac{C(W_{k-1} + (n+1-k)w_k) - C(W_{k-2} + (n+2-k)w_{k-1})}{n+1-k} \\ &\quad + \frac{C(W_t + (n-t)w_{t+1}) - C(W_{t-1} + (n+1-t)w_t)}{n-t} \\ &= IC_t + \frac{C(W_t + (n-t)w_{t+1}) - C(W_t + (n-t)w_t)}{n-t} \\ &\leq C(w_t) + \frac{C(W_t + (n-t)w_{t+1}) - C(W_t + (n-t)w_t)}{n-t}. \end{aligned}$$

If $w_{t+1} = w_t$ then $IC_{t+1} = IC_t \leq C(w_t) = C(w_{t+1})$. Now assume $w_{t+1} > w_t$. In that case, $W_t + (n-t)w_{t+1} > W_t + (n-t)w_t > w_t$. Since $C(\cdot)$ 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

$$\begin{aligned} IC_i &\leq C(w_t) + (w_{t+1} - w_t) \frac{C(w_{t+1}) - C(w_t)}{w_{t+1} - w_t} \\ &= C(w_t) + C(w_{t+1}) - C(w_t) \\ &= C(w_{t+1}). \end{aligned}$$

□

4 Fully Mixed Nash Equilibria

We prove:

Theorem 4 (Existence of Fully Mixed Nash Equilibrium). *There is always a fully mixed Nash equilibrium.*

Sketch of Proof. Fix any demand vector \mathbf{w} . Consider the fully mixed assignment \mathbf{F} , with $f_i^j = \frac{1}{m}$ for each 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 selects resource j . This occurs with probability $(1 - \frac{1}{m})^{n-1}$, and it contributes $w_i (1 - \frac{1}{m})^{n-1}$ to IC_i^j .
- Fix now any integer k , where $2 \leq k \leq n$, and assume that $k-1$ agents other than i select the resource j . This occurs with probability $(\frac{1}{m})^{k-1} (1 - \frac{1}{m})^{n-k}$. There are $\binom{n-1}{k-1}$ pure assignments where exactly $k-1$ agents (besides i) select resource j , and each 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 $\frac{1}{k} \cdot (\frac{1}{m})^{k-1} (1 - \frac{1}{m})^{n-k} \cdot \left(\binom{n-1}{k-1} w_i + \binom{n-2}{k-2} W_{-i} \right)$, where $W_{-i} = \sum_{t \in \mathcal{N}: t \neq i} w_t$.

Hence, it follows that

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

Since IC_i^j is independent of j , it follows that \mathbf{F} is a fully mixed Nash equilibrium, as needed. \square

Call \mathbf{F} , with $f_i^j = \frac{1}{m}$ for each pair of an agent $i \in \mathcal{N}$ and resource $j \in \mathcal{M}$ from the proof of Theorem 4 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 5 (Uniqueness of Standard Fully Mixed Nash Equilibrium for $n = 2$). *The standard fully mixed Nash equilibrium is the unique Nash equilibrium in the case of $n = 2$ agents with nonidentical demands.*

Sketch of Proof. Fix any demand vector \mathbf{w} . Consider any arbitrary Nash equilibrium \mathbf{P} . We will prove that (necessarily) $\mathbf{P} = \mathbf{F}$, the fully mixed Nash equilibrium from Theorem 4.

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 $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$. Without loss of generality, take that $w_1 > w_2$. Consider any resource $\ell \in \mathcal{S}_2$. Clearly,

$$\text{IC}_1 = w_1 > w_1 + \frac{w_2 - w_1}{2} p_2^\ell = w_1(1 - p_2^\ell) + \frac{w_1 + w_2}{2} p_2^\ell = \text{IC}_1^\ell,$$

a contradiction to the fact that \mathbf{P} is a Nash equilibrium. So, take any resource $j \in \mathcal{S}_1 \cap \mathcal{S}_2$. Clearly,

$$\begin{aligned} \text{IC}_1 &= \text{IC}_1^j = w_1(1 - p_2^j) + \frac{w_1 + w_2}{2} p_2^j = w_1 + \frac{w_2 - w_1}{2} p_2^j, \text{ and similarly,} \\ \text{IC}_2 &= w_2 + \frac{w_1 - w_2}{2} p_1^j. \end{aligned}$$

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

It follows that $\mathcal{S}_1 = \mathcal{S}_2$. Assume that there exists a resource $k \notin \mathcal{S}_1$. Then, $\text{IC}_2^k = w_2 < \text{IC}_2$, 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, $\text{IC}_1^j = \text{IC}_1^k$, 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 $(w_2 - w_1)p_2^j = (w_2 - w_1)p_2^k$, or $p_2^j = p_2^k$. 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. \square

The assumption of nonidentical demands in Theorem 5 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.

5 The Price of Anarchy

5.1 Lower Bound

We prove:

Theorem 6. $\mathcal{PA} \geq \frac{n}{2e}$.

Sketch of Proof. Fix any demand vector \mathbf{w} . Consider the fully mixed Nash equilibrium \mathbf{F} from Theorem 4 where each pure assignment occurs with the same probability $(\frac{1}{m})^n$. 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 she selects. In these pure assignments, the resource chosen by agent 1 is the resource with maximum Resource Cost. Thus,

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

Fix now the demand vector \mathbf{w} with $w_1 = \Theta(2^n)$ and $w_i = 1$ for all agents $i \neq 1$. Then, clearly, $\frac{w_1}{W} \geq \frac{1}{2}$. Proposition 1 implies that $\text{OPT}(\mathbf{w}) = \frac{W}{n}$. Hence, $\mathcal{PA} \geq \frac{nw_1}{W} \left(\frac{m-1}{m}\right)^{n-1} \geq \frac{n}{2} \left(\frac{m-1}{m}\right)^{n-1} \geq \frac{n}{2e}$ for $m = n$, as needed. \square

5.2 Upper Bounds

We first prove an upper bound for the special case of 2 agents.

Theorem 7. *Assume that $n = 2$. Then, $\mathcal{PA} < 2 - \frac{1}{m}$.*

Sketch of 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, which, by Proposition 5, is the unique Nash equilibrium in this case.

Note that each pure assignment occurs with the same probability $\left(\frac{1}{m}\right)^2$. 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, $\text{SC}(\mathbf{w}, \mathbf{F}) = \left(\frac{1}{m}\right)^2 (m(m-1)w_1 + m\frac{w_1+w_2}{2}) = w_1 \left(1 - \frac{1}{m}\right) + \frac{w_1+w_2}{2} \frac{1}{m}$. Since $\text{OPT}(\mathbf{w}) = \frac{w_1+w_2}{2}$, it follows that $\mathcal{PA} = \frac{2w_1(1-\frac{1}{m})}{w_1+w_2} + \frac{1}{m} < 2 - \frac{2}{m} + \frac{1}{m} = 2 - \frac{1}{m}$, as needed. \square

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

Theorem 8. $\mathcal{PA} \leq \frac{w_1}{W} \cdot n$.

Sketch of 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 $\mathcal{PA} \leq \frac{w_1}{W} = \frac{w_1}{W} n$, as needed. \square

Theorem 8 immediately implies:

Corollary 1. *Assume that for all agents $k \in \mathcal{N}$, $w_k \leq c \cdot \min_i w_i$, for some constant $c > 0$. 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 for the demands. We start by defining this special class.

Definition 1 (Bounded, Independent Probability Distributions). *The class of bounded, independent probability distributions Δ 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* [7].

Proposition 2 (Hoeffding Bound [7]). *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 $k \geq 1$. Denote $\bar{X} = \frac{1}{n} \sum_{k=1}^n X_k$ and $\bar{\mu} = \mathcal{E}(\bar{X})$. Then, for any $t > 0$,*

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

Setting $t = \varepsilon \mu$, where $0 < \varepsilon \ll 1$, and $a_k = 0$ and $b_k = \delta(n)$ for each $k \geq 1$ in Proposition 2 yields:

Corollary 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 $k \geq 1$. Denote $\bar{X} = \frac{1}{n} \sum_{k=1}^n X_k$ and $\bar{\mu} = \mathcal{E}(\bar{X})$. Then, for any $t > 0$, $\Pr \{ \bar{X} \leq (1 - \varepsilon)\mu \} \leq \exp \left(\frac{-2n\varepsilon^2 \mu^2}{\delta^2(n)} \right)$.*

We are now ready to prove:

Theorem 9. *Consider the class Δ of bounded, independent probability distributions. Then:*

1. $\mathcal{DPA}_\Delta \leq \frac{\ell_\Delta}{1 - \ell_\Delta \sqrt{\frac{1}{2 \ln n}}} + n \exp \left(-\frac{n}{\ln n} \right)$;
2. $\lim_{n \rightarrow \infty} \mathcal{DPA}_\Delta \leq \ell_\Delta$.

Sketch of Proof. Fix any probability distribution $D \in \Delta$. Fix any arbitrary demand vector \mathbf{w} drawn according to D . By Theorem 8, $\mathcal{PA} \leq \frac{w_1}{W}$. We will analyze

$\mathcal{E}_D \left(\frac{w_1}{W} \right)$. Since all demands are identically distributed, linearity of expectation

implies that $\mathcal{E}_D(\widehat{W}) = \mathcal{E}_D(w_1)$. So, for any $\varepsilon > 0$,

$$\begin{aligned} \mathcal{E}_D\left(\frac{w_1}{\widehat{W}}\right) &\leq \max_{\widehat{W} \geq (1-\varepsilon)\mathcal{E}_D(\widehat{W})} \frac{w_1}{\widehat{W}} \cdot \Pr\left\{\widehat{W} \geq (1-\varepsilon)\mathcal{E}_D(\widehat{W})\right\} \\ &\quad + \max_{\widehat{W} \leq (1-\varepsilon)\mathcal{E}_D(\widehat{W})} \frac{w_1}{\widehat{W}} \cdot \Pr\left\{\widehat{W} \leq (1-\varepsilon)\mathcal{E}_D(\widehat{W})\right\} \end{aligned}$$

Since $w_1 \leq \delta_D(n)$ and by Corollary 2,

$$\begin{aligned} \mathcal{E}_D\left(\frac{w_1}{\widehat{W}}\right) &\leq \frac{\delta_D(n)}{(1-\varepsilon)\mathcal{E}_D(\widehat{W})} \cdot 1 + n \cdot \exp\left(\frac{-2n\varepsilon^2\mathcal{E}_D^2(\widehat{W})}{\delta^2(n)}\right) \\ &\leq \frac{\delta_D(n)}{(1-\varepsilon)\mathcal{E}_D(w_1)} \cdot 1 + n \cdot \exp\left(\frac{-2n\varepsilon^2\mathcal{E}_D^2(w_1)}{\delta^2(n)}\right) \\ &\leq \frac{\ell_\Delta}{1-\varepsilon} \cdot 1 + n \cdot \exp\left(\frac{-2n\varepsilon^2}{\ell_\Delta^2}\right) \end{aligned}$$

Setting $\varepsilon = \ell_\Delta \sqrt{\frac{1}{\ln n}}$,

$$\mathcal{E}_D\left(\frac{w_1}{\widehat{W}}\right) \leq \frac{\ell_\Delta}{1 - \ell_\Delta \sqrt{\frac{1}{2 \ln n}}} \cdot 1 + n \cdot \exp\left(-\frac{n}{\ln n}\right)$$

so that $\lim_{n \rightarrow \infty} \mathcal{E}_D\left(\frac{w_1}{\widehat{W}}\right) \leq \ell_\Delta$. Since $\mathcal{DP}\mathcal{A}_\Delta = \max_{D \in \Delta} (\mathcal{E}_D(\mathcal{P}\mathcal{A}))$ and D was chosen arbitrarily, both claims follow. \square

A special 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 distribution $D \in \Delta_{sym}$, each demand w_i 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 9(2) implies:

Corollary 3. *Consider the class Δ_{sym} of bounded, independent, expectation-symmetric probability distributions. Then, $\lim_{n \rightarrow \infty} \mathcal{DP}\mathcal{A}_{\Delta_{sym}} \leq 2$.*

7 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. In particular, what are bounds on the Diffuse Price of Anarchy for the Average Cost Pricing model [1] and the Serial Cost Sharing model [14,15]?

References

1. 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.
2. G. Christodoulou, E. Koutsoupias and A. Nanavati, "Coordination Mechanisms", *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, pp. 345–357, Vol. 3142, LNCS, 2004.
3. 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*, pp. 123–134, Vol. 2380, LNCS, 2002.
4. 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*, pp. 645–657, Vol. 3142, LNCS, 2004.
5. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Extreme Nash Equilibria", *Proceedings of the 8th Italian Conference on Theoretical Computer Science*, pp. 1–20, Vol. 2841, LNCS, 2003. Full version: "The Structure and Complexity of Extreme Nash Equilibria", *Theoretical Computer Science*, accepted, 2004.
6. 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.
7. W. Hoeffding, "Probability Inequalities "for Sums of Bounded Random Variables", *Journal of the American Statistical Association*, Vol. 58, pp. 13–30, 1963.
8. 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.
9. E. Koutsoupias and C. H. Papadimitriou, "Worst-case Equilibria", *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Vol. 1563, LNCS, 1999.
10. E. Koutsoupias and C. H. Papadimitriou, "Beyond Competitive Analysis", *SIAM Journal on Computing*, Vol. 30, No. 1, pp. 300–317, April/May 2000.
11. 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*, pp. 547–558, Vol. 2996, LNCS, 2004.

12. 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*, pp. 551–561, Vol. 2747, LNCS, 2003.
13. M. Mavronicolas and P. Spirakis, “The Price of Selfish Routing”, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, 2001.
14. H. Moulin and S. Shenker, “Serial Cost Sharing”, *Econometrica*, Vol. 60, No. 5, pp. 1009–1037, September 1992.
15. H. Moulin and S. Shenker, “Average Cost Pricing versus Serial Cost Sharing: An Axiomatic Comparison”, *Journal of Economic Theory*, Vol. 64, pp. 178–201, 1994.
16. J. F. Nash, “Non-cooperative Games”, *Annals of Mathematics*, Vol. 54, No. 2, pp. 286–295, 1951.
17. C. H. Papadimitriou, “Algorithms, Games and the Internet”, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 749–753, 2001.
18. C. Shapiro and H. R. Varian, *Information Rules – A Strategic Guide to the Network Economy*, Harvard Business School Press, 1998.
19. M. Spence, “Nonlinear Prices and Welfare”, *Journal of Marketing Research*, Vol. 8, pp. 1–18, 1976.
20. H. R. Varian, “Prime Discrimination and Social Welfare”, *American Economic Review*, Vol. 75, No. 4, pp. 870–875, 1985.