

Stochastic Congestion Games with Risk-Averse Players^{*}

Haris Angelidakis¹, Dimitris Fotakis², and Thanasis Lianeas²

¹ Toyota Technological Institute at Chicago, USA
hangel@ttic.edu

² Electrical and Computer Engineering, National Technical University of Athens, Greece
fotakis@cs.ntua.gr, tlianeas@mail.ntua.gr

Abstract. Congestion games ignore the stochastic nature of resource delays and the risk-averse attitude of the players to uncertainty. To take these aspects into account, we introduce two variants of atomic congestion games, one with *stochastic players*, where each player assigns load to her strategy independently with a given probability, and another with *stochastic edges*, where the latency functions are random. In both variants, the players are risk-averse, and their individual cost is a player-specific quantile of their delay distribution. We focus on parallel-link networks and investigate how the main properties of such games depend on the risk attitude and on the participation probabilities of the players. In a nutshell, we prove that stochastic congestion games on parallel-links admit an efficiently computable pure Nash equilibrium if the players have either the same risk attitude or the same participation probabilities, and also admit a potential function if the players have the same risk attitude. On the negative side, we present examples of stochastic games with players of different risk attitudes that do not admit a potential function. As for the inefficiency of equilibria, for parallel-link networks with linear delays, we prove that the Price of Anarchy is $\Theta(n)$, where n is the number of stochastic players, and may be unbounded, in case of stochastic edges.

1 Introduction

Congestion games provide an elegant and useful model of selfish resource allocation in large-scale networks. In an (atomic) *congestion game*, a finite set of players, each with an unsplittable unit of load, compete over a finite set of resources (or edges). All players using an edge experience a latency given by a non-negative and non-decreasing function of the edge's load (or congestion). Each player selects a path between her origin and destination, trying to minimize her *individual cost*, that is, the sum of the latencies on the edges in the chosen path. A natural solution concept is that of a *pure Nash equilibrium* (PNE), a configuration where no player can decrease her individual cost by unilaterally changing her path.

In a seminal work, Rosenthal [18] proved that the PNE of congestion games correspond to the local optima of a natural potential function, and thus every congestion

^{*} This research was supported by the project Algorithmic Game Theory, co-financed by the European Union (European Social Fund - ESF) and Greek national funds, through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: THALES, investing in knowledge society through the European Social Fund.

game admits a PNE. Following [18], the properties of congestion games and several variants of them have been extensively studied. The prevailing questions in recent work have to do with whether congestion games and some natural generalizations of them admit an (approximate) PNE and/or an (approximate) potential function (see e.g., [13], [11], [12] and [5]), with bounding the convergence time to a PNE if the players act selfishly (see e.g., [1], [7], and [10]), and with quantifying the inefficiency of PNE due to the players' selfish behavior (see e.g., [4], [8], [2], and [6]). Notably, a significant part of recent work concerns the properties of congestion games and their generalizations on parallel-link networks (see e.g., [13], [6], and [10], and the references therein).

However, most research work on congestion games essentially ignores the stochastic nature of edge delays and assumes that players have precise knowledge of the (deterministic) edge latencies. On the contrary, in real life situations, players cannot accurately predict the actual edge delays, not only because they cannot know the exact congestion of every edge, but also due to (a priori unknown) external events (e.g., some construction work, a minor accident, a link failure) that may affect the edge latencies and introduce uncertainty. It is therefore natural to assume that the players decide on their strategies based only on estimations of their actual delay and, most importantly, that they are fully aware of the uncertainty and of the potential inaccuracy of their estimations. So, to secure themselves from the event of an increased delay, players select their paths taking uncertainty into account (e.g., people either take a safe route or plan for a larger than usual delay when they head to an important meeting).

Such considerations give rise to congestion games with stochastic delays and risk-averse players, where instead of the path that minimizes her expected delay, each player selects a path that guarantees her a reasonably low actual delay with reasonably high confidence. Here, the actual delay of each player can be modeled by a random variable. Then, a common assumption is that players seek to minimize either a convex combination of the expectation and the variance of their delay, or a player-specific quantile of the delay distribution (see also [19], [9] about the cost functions of risk-averse players, and [17] about possible ways of risk quantification in optimization under uncertainty).

Previous Work. Following the research direction above, Ordóñez and Stier-Moses [15] considered *nonatomic* congestion games and suggested that each path should be penalized by an additive term that increases with the risk-aversion of the players and with the maximum deviation from the expected delay of the path (however, this term does not depend on the actual load of the edges). For each path, the additive term can be chosen either as a δ -fraction of (resp. a δ -quantile of a random variable depending on) the maximum deviation from the expected delay of the path, or simply, as the sum of the δ -fractions of the maximum deviation from the expected delay of each edge in the path, where δ quantifies the risk-aversion of the players. Under some general assumptions, [15] proves that an equilibrium exists and is essentially unique in all the cases above.

Subsequently, Nikolova and Stier-Moses [14] suggested a model of stochastic selfish routing with risk-averse players, where each player selects a path that minimizes the expected delay plus δ times the standard deviation of the delay, where δ quantifies the risk-aversion of the players. They considered nonatomic and atomic congestion games, mostly with homogeneous players, that share the same risk attitude, and distinguished between the case where the standard deviation of a path's delay is *exogenous*, i.e., it does

not depend on the load of the edges in the path, and the case where it is *endogenous*, i.e., it is a function of the load. They proved that in the exogenous case, which is similar to the model of [15], stochastic routing games essentially retain the nice properties of standard congestion games: they admit a potential function and, in the nonatomic setting, a unique equilibrium, and the inefficiency of equilibria can be bounded as for standard congestion games. In the endogenous case, they proved that nonatomic stochastic routing games admit an equilibrium, which is not necessarily unique, but may not admit a cardinal potential. Moreover, atomic stochastic routing games may not admit a PNE even in simple extension-parallel networks with 2 players and linear delays.

Contribution. Following this research agenda, we seek a better understanding of the properties of congestion games with stochastic delays and risk-averse players. We focus on atomic congestion games and introduce two variants of stochastic congestion games. We start from the observation that the variability of edge delays comes either from the variability of the traffic demand, and the subsequent variability of the edge loads, or from the variability of the edge performance level. Decoupling them, we introduce two variants, namely *Congestion Games with Stochastic Players* and *Congestion Games with Stochastic Edges*, each capturing one of the two sources of uncertainty above.

Congestion Games with Stochastic Players aim to model the variability of the traffic demand. Specifically, each player i participates in the game independently with probability p_i . As a result, the total network load, the edge loads, and the edge and the path latencies are all random variables. On the other hand, Congestion Games with Stochastic Edges aim to model variability in the network operation. Now, each edge e may operate either at the “standard” mode, where its latency is given by a function $f_e(x)$, or at the “faulty” mode (e.g., after a minor accident or a link failure), where its latency is given by $g_e(x)$, with $g_e(x) \geq f_e(x)$. Each edge e switches to the “faulty” mode independently with a given probability p_e . Hence, the network load and the edge loads are now deterministic, but the edge and the path latencies are random variables. In both variants, players are risk-averse to the stochastic delays. Specifically, each player i has a (possibly different) desired *confidence level* δ_i , and her cost on a path q is the δ_i -quantile (a.k.a. value-at-risk) of the delay distribution of q . In words, the individual cost of player i is the minimum delay she can achieve along q with probability at least δ_i .

At the conceptual level, the model of Congestion Games with Stochastic Players is similar to the model with endogenous standard deviations of [14]. In fact, using Chernoff bounds, one can show that for linear latency functions, if the expected edge loads are not too small, our δ_i -quantile individual cost can be approximated by the individual cost used in [14]. However, we also consider stochastic demands, a direction suggested in [14, Sec. 7] to enrich the model, and players that are heterogeneous with respect to risk attitude. As for Congestion Games with Stochastic Edges, the model is conceptually similar to the model with exogenous standard deviations of [14].

In the technical part of the paper, we restrict ourselves to parallel-link networks with symmetric player strategies, and investigate how the properties of stochastic congestion games depend on the players’ participation probabilities and confidence levels. We first observe that such games admit a potential function and an efficiently computable PNE, if the players are homogeneous, namely if they have the same confidence level δ and, in case of stochastic players, the same participation probability p (Theorems 1 and 7).

We also show that if the players have different confidence levels (and the same participation probability, if they are stochastic), stochastic congestion games belong to the class of player-specific congestion games [13], and thus admit a PNE computable in polynomial time (Corollaries 1 and 2). On the negative side, we prove that such games may not admit a potential function (Theorems 2 and 8). For Congestion Games with Stochastic Players that have the same confidence level and different participation probabilities, we show that they admit a lexicographic potential (Theorem 4), and thus a PNE, which can be computed by a simple greedy best response algorithm (Theorem 3). As for the inefficiency of PNE, in the case of linear latency functions, we prove that the Price of Anarchy (PoA) is $\Theta(n)$, if we have n stochastic players, and (Theorems 5 and 6), and may be unbounded, in the case of stochastic edges (Theorem 9).

Other Related Work. There is a significant volume of work on theoretical and practical aspects of transportation networks with uncertain delays, which however focuses on nonatomic games and adopts notions of individual cost and viewpoints quite different from ours (see e.g., the discussion in [14]). Motivated by applications with only partial knowledge of the number of players participating in the game, Ashlagi, Monderer, and Tennenholtz [3] considered congestion games on parallel links with stochastic players. However, the players in [3] are risk-neutral, since their individual cost is the expected delay of the chosen link. They proved that a generalization of the fully mixed equilibrium remains a mixed Nash equilibrium in this setting. Very recently, Piliouras, Nikolova, and Shamma [16] considered atomic congestion games with risk-averse players and delays determined by a randomized scheduler of the players on each edge. They obtained tight bounds on the PoA of such games with linear latencies under various notions of risk-averse individual cost. Interestingly, they proved that the PoA can be unbounded for the individual cost of [14].

2 Notation and Preliminaries

In this section, we introduce the notation and the basic terminology of standard congestion games. For a random variable X , we let $\mathbb{E}[X]$ denote the *expectation* and $\mathbb{V}\text{ar}[X]$ denote the *variance* of X . For an event E , we let $\mathbb{P}\text{r}[E]$ denote its probability. For a vector $x = (x_1, \dots, x_n)$, we let $x_{-i} \equiv (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i}, x'_i) \equiv (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

Congestion Games. A *congestion game* is a tuple $\mathcal{G}(N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E})$, where N is the set of players, E is the set of resources, $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$ is the strategy space of each player i , and $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ is a non-negative and non-decreasing latency function associated with each resource e . A congestion game is *symmetric* if all players share the same strategy space. In what follows, we let n denote the number of players and m denote the number of resources.

A *configuration* is a vector $\sigma = (\sigma_1, \dots, \sigma_n)$ consisting of a strategy $\sigma_i \in \Sigma_i$ for each player i . For every resource e , we let $\sigma_e = |\{i \in N : e \in \sigma_i\}|$ denote the congestion induced on e by σ . The individual cost of player i in the configuration σ is $c_i(\sigma) = \sum_{e \in \sigma_i} d_e(\sigma_e)$. A configuration σ is a *pure Nash equilibrium* (PNE) if no player can improve her individual cost by unilaterally changing her strategy. Formally, σ is a PNE if for every player i and every strategy $s_i \in \Sigma_i$, $c_i(\sigma) \leq c_i(\sigma_{-i}, s_i)$.

Next, we focus on symmetric congestion games on parallel-link networks, where the strategies are singletons and there is a strategy for every resource. Thus, we use the terms “resource” and “edge”, and “strategy” and “path” interchangeably.

Social Cost. To quantify the inefficiency of PNE, configurations are usually evaluated using the *total cost* of the players. In standard congestion games, the total cost of a configuration σ , denoted $C(\sigma)$, is $C(\sigma) = \sum_{i \in N} c_i(\sigma)$. The *optimal configuration*, usually denoted o , minimizes the total cost among all possible configurations.

Price of Anarchy. The (pure) *Price of Anarchy* (PoA) of a congestion game \mathcal{G} is the maximum ratio $C(\sigma)/C(o)$ over all PNE σ of \mathcal{G} . The PoA of a class of games is defined as the maximum PoA among all games in the class.

3 Congestion Games with Stochastic Players

3.1 The Model

In *Congestion Games with Stochastic Players*, each player i is described by a tuple (p_i, δ_i) , where $p_i \in [0, 1]$ is the probability that player i participates in the game, by assigning a unit of load to her strategy, and $\delta_i \in [\frac{1}{2}, 1]$ is the *confidence level* (or risk-aversion) of player i . Essentially, each player i is associated with a Bernoulli random variable X_i that is 1 with probability p_i , and 0 with probability $1 - p_i$. Then, the load of each edge e in a configuration σ is the random variable $N_e(\sigma) = \sum_{i: e \in \sigma_i} X_i$, and the cost of a strategy q in σ is the random variable $D_q(\sigma) = \sum_{e \in q} d_e(N_e(\sigma))$.

Given that player i participates in the game, the delay of player i in σ is given by the random variable:

$$D_i(\sigma) = \sum_{e \in \sigma_i} d_e \left(1 + \sum_{j \neq i: e \in \sigma_j} X_j \right).$$

We note that conditional on $X_i = 1$, $D_i(\sigma) = D_{\sigma_i}(\sigma)$, i.e., the delay of i in σ is equal to the cost of her strategy in σ , conditional that i participates in the game.

The (risk-averse) individual cost $c_i(\sigma)$ perceived by player i in σ is the δ_i -quantile (or value-at-risk) of $D_i(\sigma)$. Formally, $c_i(\sigma) = \min\{t : \Pr[D_i(\sigma) \leq t] \geq \delta_i\}$. We note that for parallel-link networks, the (risk-averse) individual cost of the players can be computed efficiently. PNE are defined as before, but with respect to the risk-averse individual cost of the players.

Depending on whether players have the same participation probabilities p_i and/or the same confidence levels δ_i , we distinguish between four classes of Congestion Games with Stochastic Players:

- *homogeneous*, where all players have the same participation probability p and confidence level δ .
- *p-homogeneous*, where all players have the same participation probability p , but may have different confidence levels.
- *δ -homogeneous*, where all players have the same confidence level δ , but may have different participation probabilities.
- *heterogeneous*, where both the participation probabilities and the confidence levels may be different.

3.2 Stochastic Players on Parallel Links: Existence and Computation of PNE

In the following, we restrict ourselves to parallel-link networks, and investigate the existence and the efficient computation of PNE for the four cases considered above.

Homogeneous Stochastic Players. If the players are homogeneous, stochastic congestion games on parallel-links are equivalent to standard congestion games on parallel-links (but with possibly different latencies), because the (risk-averse) individual cost of each player in a configuration σ depends only on the link e and its congestion σ_e . The proof of the following employs Rosenthal's potential function and a Greedy Best Response dynamics that guarantee the existence and the efficient computation of a PNE.

Theorem 1. *Congestion Games with Homogeneous Stochastic Players on parallel-link networks admit an exact potential function. Moreover, a PNE can be computed in polynomial time.*

p -Homogeneous Stochastic Players. In this case, a stochastic game is equivalent to a congestion game on parallel links with player-specific costs [13], as the (risk-averse) individual cost of each player i in a configuration σ depends only on the link e , its congestion σ_e , and i 's confidence level δ_i . Thus, we obtain that:

Corollary 1. *Congestion Games with p -Homogeneous Stochastic Players on parallel-link networks admit a PNE. Moreover, a PNE can be computed in polynomial time.*

Milchtaich [13] proved that parallel-link games with general player-specific costs may not admit a potential function. However, in our case the players' individual costs are correlated, as for any edge, there is a common distribution on which they depend. Nevertheless, we next show that parallel-link games with p -homogeneous stochastic players and linear latencies may not admit any (even generalized) potential function.

Theorem 2. *There are Congestion Games with p -Homogeneous Stochastic Players on parallel-link networks with linear delays that do not admit any potential function.*

Proof. It suffices to show that there is an infinite sequence of deviations in which each deviating player improves her cost. To this end, we adjust the example in [13, Section 5] to our setting. We recall that since players have the same participation probability p , the load on each edge e that player i considers is binomially distributed.

We let $p = 0.75$, and consider 3 parallel links, e_1 , e_2 , and e_3 , 3 "special" players, that change their strategies and form a better response cycle, with $\delta_1 = 0.75$, $\delta_2 = 0.58$ and $\delta_3 = 0.6$, and $n_1 = 25$ additional players on e_1 , $n_2 = 20$ additional players on e_2 and $n_3 = 9$ additional players on e_3 . The latency functions of the 3 edges are $f_1(k) = 3k + 71$, $f_2(k) = 6k + 33$ and $f_3(k) = 15k + 1$.

We proceed to describe a better response cycle that consists of 6 different configurations $\sigma_1, \dots, \sigma_6$. Each configuration is represented by a vector $[S_1, S_2, S_3]$, where S_i is the subset of the "special" players using edge e_i .

$$\begin{aligned} \sigma_1 &= [\{1, 2\}, \{3\}, \emptyset] \rightarrow \sigma_2 = [\{1, 2\}, \emptyset, \{3\}] \rightarrow \sigma_3 = [\{2\}, \emptyset, \{1, 3\}] \rightarrow \\ \sigma_4 &= [\emptyset, \{2\}, \{1, 3\}] \rightarrow \sigma_5 = [\emptyset, \{2, 3\}, \{1\}] \rightarrow \sigma_6 = [\{1\}, \{2, 3\}, \emptyset] \rightarrow \sigma_1 \end{aligned}$$

To verify that this is indeed a better response cycle, we give the vectors of the risk-averse individual cost of the “special” players in each configuration: $c(\sigma_1) = (137, 134, 135)$, $c(\sigma_2) = (137, 134, 121)$, $c(\sigma_3) = (136, 131, 136)$, $c(\sigma_4) = (136, 129, 136)$, $c(\sigma_5) = (136, 135, 135)$, and $c(\sigma_6) = (134, 135, 135)$. \square

δ -Homogeneous Stochastic Players. In this case, players have the same confidence level δ , but their participation probabilities may be different. We next show how to efficiently compute a PNE in parallel-link networks by the p -Decreasing Greedy Best Response algorithm, or p -DGBR, in short, which proceeds as follows:

- Order the players in *non-increasing* order of their participation probabilities p_i .
- Assign the current player, in the previous order, to the edge corresponding to her best response strategy in the current configuration.
- Repeat until all players are added.

Theorem 3. *The p -DGBR algorithm computes, in $O(nm + n^2)$ time, a PNE for Congestion Games with δ -Homogeneous Stochastic Players on parallel-link networks with general latency functions.*

Proof. The proof is by induction on the number of players. We assume that we are at a PNE, and player i is assigned to edge e . Since players on other edges do not deviate, we only have to show that players on e do not deviate. Let k be any player already on e , which implies that $p_k \geq p_i$. It suffices to show that in the current configuration σ , with $\sigma_i = \sigma_k = e$, we have that $c_k(\sigma) \leq c_i(\sigma)$. This holds because $p_k \geq p_i$ and players i and k perceive the same cost on any other edge.

Formally, let us consider $c_k(\sigma)$ and $c_i(\sigma)$. We have that:

$$c_k(\sigma) = \min \left\{ t : \Pr \left[d_e \left(1 + X_i + \sum_{j \neq i, k: \sigma_j = e} X_j \right) \leq t \right] \geq \delta \right\} \text{ and}$$

$$c_i(\sigma) = \min \left\{ t : \Pr \left[d_e \left(1 + X_k + \sum_{j \neq i, k: \sigma_j = e} X_j \right) \leq t \right] \geq \delta \right\}.$$

Since $p_k \geq p_i$, for any $r \in \mathbb{N}$, we have that:

$$\begin{aligned} \Pr \left[X_k + \sum_{j \neq i, k: \sigma_j = e} X_j \leq r \right] &= \Pr \left[\sum_{j \neq i, k: \sigma_j = e} X_j \leq r \right] - \Pr \left[\sum_{j \neq i, k: \sigma_j = e} X_j = r \right] p_k \\ &\leq \Pr \left[\sum_{j \neq i, k: \sigma_j = e} X_j \leq r \right] - \Pr \left[\sum_{j \neq i, k: \sigma_j = e} X_j = r \right] p_i \\ &= \Pr \left[X_i + \sum_{j \neq i, k: \sigma_j = e} X_j \leq r \right]. \end{aligned}$$

Thus, since the edge latency functions are non-decreasing, we obtain that:

$$\begin{aligned} \Pr \left[d_e \left(1 + X_k + \sum_{j \neq i, k: \sigma_j = e} X_j \right) \leq d_e(r + 1) \right] \\ \leq \Pr \left[d_e \left(1 + X_i + \sum_{j \neq i, k: \sigma_j = e} X_j \right) \leq d_e(r + 1) \right]. \end{aligned}$$

Therefore, $c_k(\sigma) \leq c_i(\sigma)$, as required. The total computation time is $O(nm + n^2)$, as at each step i , the computations for the newly inserted player take $O(m + i^2)$ time, and we can use memoization to avoid recalculations. \square

We next show that Congestion Games with δ -Homogeneous Stochastic Players admit a two-dimensional lexicographic potential function.

Theorem 4. *Congestion Games with δ -Homogeneous Stochastic Players on parallel-link networks admit a generalized potential function.*

Proof. We define, for each edge e and each configuration σ , a two-dimensional vector $v_{e,\sigma}$ and a total order on these vectors. Moreover, for each configuration σ , we define a vector $w_\sigma = (v_{e,\sigma})_{e \in E}$, where the vectors $v_{e,\sigma}$ appear in increasing lexicographic order. The crux of the proof is to show that for any improving deviation that changes the configuration from σ to σ' , we have that $w_\sigma < w_{\sigma'}$. Thus, any decreasing function on the vectors w_σ can serve as a generalized potential function.

Formally, we let $c_e(\sigma) = \min\{t : \Pr[d_e(1 + N_e(\sigma)) \leq t] \geq \delta\}$ be the *outside δ -cost* of each edge e under σ , i.e. the cost that any player not in e perceives when she considers moving to e . By definition, we have that:

$$c_e(\sigma) = c_i(\sigma_{-i}, e) \quad \forall i : \sigma_i \neq e \quad (1)$$

$$c_e(\sigma) \geq c_i(\sigma) \quad \forall i : \sigma_i = e \quad (2)$$

We let $v_{e,\sigma} = (c_e(\sigma), \sigma_e)$, and consider the lexicographic order on these pairs:

- $(x_1, y_1) < (x_2, y_2)$, if either $x_1 < x_2$ or $x_1 = x_2$ and $y_1 < y_2$.
- $(x_1, y_1) = (x_2, y_2)$, if $x_1 = x_2$ and $y_1 = y_2$.
- $(x_1, y_1) > (x_2, y_2)$, otherwise.

For any configuration σ , we let $w_\sigma = (v_{e,\sigma})_{e \in E}$ be the vector consisting of the pairs $v_{e,\sigma}$ in increasing lexicographic order. We next show that after any improving deviation, the new configuration σ' has $w_{\sigma'} > w_\sigma$.

Let us assume that player i performs an improving deviation from e to e' , and let $\sigma = (\sigma_{-i}, e)$ be the initial configuration and $\sigma' = (\sigma_{-i}, e')$ be the final configuration. Since we consider an improving deviation of player i , $c_i(\sigma) > c_i(\sigma')$. Furthermore, by (1), $c_i(\sigma') = c_i(\sigma_{-i}, e') = c_{e'}(\sigma)$, and by (2), $c_e(\sigma) \geq c_i(\sigma)$. Thus, we obtain that $c_e(\sigma) > c_{e'}(\sigma)$, which implies that $v_{e',\sigma} < v_{e,\sigma}$. Hence, if we consider the coordinates of w_σ , we have that the pair $v_{e',\sigma}$ of e' appears before the pair $v_{e,\sigma}$ of e .

Since we consider a deviation from e to e' , the only pairs affected are $v_{e,\sigma}$ and $v_{e',\sigma}$. Consequently, in order to show that $w_\sigma < w_{\sigma'}$, we need to show (i) that $v_{e',\sigma} < v_{e',\sigma'}$, and (ii) that $v_{e',\sigma} < v_{e,\sigma'}$. In words, we need to show that the pair of e' increases by i 's move from e to e' , and that the pair of e in σ' , although it might decrease, still remains greater than the pair of e' in σ .

As for inequality (i), we observe that $\sigma_{e'} < \sigma'_{e'}$ and that $c_{e'}(\sigma) \leq c_{e'}(\sigma')$. Combining these inequalities, we conclude that $v_{e',\sigma} < v_{e',\sigma'}$.

To show inequality (ii), we combine (1) with the hypothesis that player i performs an improving deviation from e to e' , and obtain that $c_i(\sigma) > c_i(\sigma_{-i}, e') = c_{e'}(\sigma)$. Also, considering the outside δ -cost of e in σ' and using that $\sigma'_i \neq e$, we obtain that

$c_e(\sigma') = c_i(\sigma'_{-i}, e) = c_i(\sigma)$, because σ'_{-i} and σ_{-i} are identical. Combining these, we conclude that $c_{e'}(\sigma) < c_e(\sigma')$, which immediately implies that $v_{e',\sigma} < v_{e,\sigma'}$.

We have thus established a correspondence between configurations σ and the vectors w_σ , and that for any improving deviation that changes the configuration from σ to σ' , $w_\sigma < w_{\sigma'}$. Now, let us consider any strictly decreasing function Φ from the vectors w_σ to \mathbb{R} . Then, for any configuration σ , any edges e, e' , and any player i , we have that

$$c_i(\sigma_{-i}, e) > c_i(\sigma_{-i}, e') \Rightarrow \Phi(\sigma_{-i}, e) > \Phi(\sigma_{-i}, e').$$

Consequently, Φ serves as generalized potential function for Congestion Games with δ -Homogeneous Stochastic Players. \square

3.3 The Price of Anarchy for Stochastic Games with Affine Latencies

In Congestion Games with Stochastic Players, we let the *total cost* of a configuration σ be $C(\sigma) = \mathbb{E}\left[\sum_{i \in N} X_i D_i(\sigma)\right]$, which is a natural generalization of the total cost for standard congestion games. We let o denote an optimal configuration that minimizes the total cost. Then, as for standard congestion games, the Price of Anarchy (PoA) of a stochastic congestion game \mathcal{G} is the maximum ratio $C(\sigma)/C(o)$ over all PNE σ of \mathcal{G} .

Next, we first convert the total cost $C(\sigma)$ to a more convenient form, and then present upper and lower bounds on the PoA of Stochastic Congestion Games with Stochastic Players and affine latency functions.

As observed in Section 3.1, if we condition on $X_i = 1$, i.e., that player i participates in the game, $D_i(\sigma) = D_{\sigma_i}(\sigma)$, and thus, $X_i D_i(\sigma) = X_i D_{\sigma_i}(\sigma)$. Therefore,

$$C(\sigma) = \mathbb{E}\left[\sum_{i \in N} X_i D_{\sigma_i}(\sigma)\right] = \sum_e \mathbb{E}[N_e(\sigma) d_e(N_e(\sigma))].$$

Hence, for affine latency functions $d_e(x) = a_e x + b_e$, we have that

$$\begin{aligned} C(\sigma) &= \sum_e \mathbb{E}[N_e(\sigma)(a_e N_e(\sigma) + b_e)] \\ &= \sum_e [a_e(\mathbb{E}[N_e(\sigma)]^2 + \text{Var}[N_e(\sigma)]) + b_e \mathbb{E}[N_e(\sigma)]]. \end{aligned}$$

Theorem 5. *Congestion Games with n Stochastic Players on parallel-link networks with affine latency functions have $\text{PoA} = O(n)$.*

Proof. Let $d_e(x) = a_e x + b_e$ denote the affine latency of each edge e . We first observe that (i) since $\delta \geq 1/2$, the cost that a player i perceives on her edge e is at least as large as her expected delay on e due to the load caused by the other players on e , formally $c_i(\sigma) \geq a_e \mathbb{E}[\sum_{j \neq i: e=\sigma_j} X_j] + b_e$, and that (ii) at equilibrium, all players perceive a cost of at most $n(a + b)$, where $a + b = \min_e \{a_e + b_e\}$, since otherwise, some player would have an incentive to deviate to the edge with latency $a x + b$.

In what follows, we let f be any PNE, and let o be an optimal configuration. Based on the observations above, we next show that $C(f) \leq 3nC(o)$.

For convenience, we let $F_e = N_e(f)$ and $O_e = N_e(o)$. We have that:

$$\begin{aligned} C(f) &= \sum_e \left(a_e (\mathbb{E}[F_e]^2 + \text{Var}[F_e]) + b_e \mathbb{E}[F_e] \right) \\ &= \sum_e \mathbb{E}[F_e] \left(a_e \mathbb{E}[F_e] + b_e + a_e \frac{\text{Var}[F_e]}{\mathbb{E}[F_e]} \right) \\ &\leq \sum_e \mathbb{E}[F_e] \left(c_{\max} + a_e + a_e \frac{\text{Var}[F_e]}{\mathbb{E}[F_e]} \right) \leq 3 \sum_e \mathbb{E}[F_e] c_{\max}, \end{aligned}$$

where c_{\max} denotes the largest cost of a player in f . The inequalities follow from observation (i) above, from $c_{\max} \geq a_e$ for all used edges e , and from $\text{Var}[F_e] \leq \mathbb{E}[F_e]$.

Using observation (ii) above, with $a + b = \min_e \{a_e + b_e\}$, we obtain that:

$$\begin{aligned} C(f) &\leq 3 c_{\max} \sum_e \mathbb{E}[F_e] \leq 3n(a + b) \sum_e \mathbb{E}[F_e] = 3n(a + b) \sum_{i \in N} p_i \\ &= 3n(a + b) \sum_e \mathbb{E}[O_e] \leq 3n \sum_e \mathbb{E}[O_e] (a_e + b_e) \\ &\leq 3n \sum_e \mathbb{E}[O_e] \left(a_e \frac{\mathbb{E}[O_e]^2 + \text{Var}[O_e]}{\mathbb{E}[O_e]} + b_e \right) = 3nC(o), \end{aligned}$$

where the last inequality follows from $\mathbb{E}[O_e]^2 + \text{Var}[O_e] = \mathbb{E}[O_e^2] \geq \mathbb{E}[O_e]$. \square

Theorem 6. *There are Congestion Games with n Homogeneous Stochastic Players on parallel-link networks with affine latency functions that have $\text{PoA} = \Omega(n)$.*

Proof sketch. We consider a game with n stochastic players on $k + 1$ parallel edges. Edge e_1 has latency $d_1(x) = x$, and every other edge e_j has latency $d_j(x) = (n - k)x$, $j = 2, \dots, k + 1$. The players have participation probability p and confidence level $\delta = 1$. The configuration where $n - k$ players use e_1 and each of the remaining k players uses a different edge e_j , $j = 2, \dots, k + 1$, is a PNE. In the optimal configuration, all n players are assigned to e_1 . Calculating the total cost of these configurations, and using $k = n/2$ and $p = 1/n$, we obtain that the PoA is roughly $n/8$. \square

4 Congestion Games with Stochastic Edges

The Model. In *Congestion Games with Stochastic Edges*, players are deterministic, i.e., they always participate in the game. As before, each player i has a confidence level $\delta_i \in [\frac{1}{2}, 1]$. On the other hand, edges have a stochastic behavior, in the sense that the latency function of each edge e is an independent random variable:

$$d_e(x) = \begin{cases} f_e(x) & \text{with probability } 1 - p_e \\ g_e(x) & \text{with probability } p_e. \end{cases}$$

The delay of edge e under congestion k is given by the random variable $X_e(k)$, which is equal to $f_e(k)$, with probability $1 - p_e$, and to $g_e(k)$, with probability p_e , and the delay of

a player i in a configuration σ is given by the random variable $D_i(\sigma) = \sum_{e \in \sigma_i} X_e(\sigma_e)$. The risk-averse individual cost of player i in σ is $c_i(\sigma) = \min\{t : \Pr[D_i(\sigma) \leq t] \geq \delta_i\}$, and the total cost of σ is $C(\sigma) = \mathbb{E}\left[\sum_{i \in N} D_i(\sigma)\right]$.

For congestion Games with Stochastic Edges, we distinguish between the case of *homogeneous* players, where all players have the same confidence level δ , and the case of *heterogeneous* players, where each player i may have a different confidence level δ_i .

4.1 Stochastic Edges on Parallel Links: Existence and Computation of PNE

Next, we restrict ourselves to Congestion Games with Stochastic Edges on parallel-link networks, and investigate the existence and the efficient computation of PNE.

Homogeneous Players. If the players are homogeneous, any Congestion Game on stochastic parallel links can be transformed into a standard congestion game on parallel links, but possibly with different latency functions. This holds because the risk-averse individual cost of each player in a configuration σ depends only on the link e and its congestion σ_e . Based on this observation, we can show that:

Theorem 7. *Stochastic Congestion Games with Stochastic Edges and Homogeneous Players on parallel-link networks admit an exact potential function. Moreover, a PNE can be computed in $O(nm)$ time.*

Heterogeneous Players. In this case, a Congestion Game on stochastic parallel links is a congestion game on parallel links with player-specific costs [13]. This holds because the risk-averse individual cost of each player i in a configuration σ depends only on the link e , its congestion σ_e , and i 's confidence level δ_i . Thus, we obtain that:

Corollary 2. *Congestion Games with Stochastic Edges and Heterogeneous Players on parallel-link networks admit a PNE computable in polynomial time.*

Milchtaich [13] proved that parallel-link games with general player-specific costs may not admit a potential function. But here, as in Section 3.2, the players' individual costs on each edge are correlated with each other. Nevertheless, the following shows that Congestion Games with Stochastic Edges do not admit any (even generalized) potential function.

Theorem 8. *There are Congestion Games with Stochastic Edges and Heterogeneous Players on parallel-link networks with affine latency functions that do not admit any potential function.*

4.2 Price of Anarchy

The following shows that selfish risk-averse players on stochastic parallel links may cause an unbounded degradation in the network performance at equilibrium.

Theorem 9. *There are Congestion Games with Stochastic Edges and Homogeneous Players on parallel-link networks with affine latencies that have an unbounded PoA.*

Acknowledgements. We wish to thank Christos Tzamos for his help in the lexicographic ordering argument, used in the proof of Theorem 4.

References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the Impact of Combinatorial Structure on Congestion Games. *Journal of the Association for Computing Machinery* 55(6), 25 (2008)
2. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In: Durand, B., Thomas, W. (eds.) *STACS 2006*. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006)
3. Ashlagi, I., Monderer, D., Tennenholtz, M.: Resource selection games with unknown number of players. In: *Proc. of the 5th Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pp. 819–825 (2006)
4. Awerbuch, B., Azar, Y., Epstein, A.: The Price of Routing Unsplittable Flow. In: *Proc. of the 37th ACM Symposium on Theory of Computing (STOC 2005)*, pp. 57–66 (2005)
5. Caragiannis, I., Fanelli, A., Gravin, N., Skopalik, A.: Efficient computation of approximate pure Nash equilibria in congestion games. In: *Proc. of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS 2011)*, pp. 532–541 (2011)
6. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight Bounds for Selfish and Greedy Load Balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006*. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
7. Chien, S., Sinclair, A.: Convergence to Approximate Nash Equilibria in Congestion Games. In: *Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, pp. 169–178 (2007)
8. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: *Proc. of the 37th ACM Symposium on Theory of Computing (STOC 2005)*, pp. 67–73 (2005)
9. Fiat, A., Papadimitriou, C.: When the players are not expectation maximizers. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) *SAGT 2010*. LNCS, vol. 6386, pp. 1–14. Springer, Heidelberg (2010)
10. Fotakis, D.: Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory of Computing Systems* 47(1), 113–136 (2010)
11. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish Unsplittable Flows. *Theoretical Computer Science* 348, 226–239 (2005)
12. Harks, T., Klimm, M.: On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research* 37(3), 419–436 (2012)
13. Milchtaich, I.: Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior* 13, 111–124 (1996)
14. Nikolova, E., Stier-Moses, N.E.: Stochastic selfish routing. In: Persiano, G. (ed.) *SAGT 2011*. LNCS, vol. 6982, pp. 314–325. Springer, Heidelberg (2011)
15. Ordóñez, F., Stier-Moses, N.: Wardrop equilibria with risk-averse users. *Transportation Science* 44(1), 63–86 (2010)
16. Piliouras, G., Nikolova, E., Shamma, J.S.: Risk sensitivity of price of anarchy under uncertainty. In: *Proc. of the 14th ACM Conference on Electronic Commerce (EC 2013)*, pp. 715–732 (2013)
17. Rockafellar, R.T.: Coherent Approaches to Risk in Optimization Under Uncertainty. In: *Tutorials in Operations Research*, pp. 38–61 (2007)
18. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
19. Tversky, A., Kahneman, D.: Prospect Theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291 (1979)