

Congestion Games with Agent Failures

Reshef Meir¹ and Moshe Tennenholtz² and Yoram Bachrach³ and Peter Key³

¹The Hebrew University of Jerusalem, Israel and Microsoft Research, Herzlia
reshef.meir@mail.huji.ac.il

²Technion-Israel Institute of Technology and Microsoft Research, Herzlia
moshet@microsoft.com

³Microsoft Research, Cambridge UK
yorambac@gmail.com, peter.key@microsoft.com

Abstract

We propose a natural model for agent failures in congestion games. In our model, each of the agents may fail to participate in the game, introducing uncertainty regarding the set of active agents. We examine how such uncertainty may change the Nash equilibria (NE) of the game. We prove that although the perturbed game induced by the failure model is not always a congestion game, it still admits at least one pure Nash equilibrium. Then, we turn to examine the effect of failures on the maximal social cost in any NE of the perturbed game.

We show that in the limit case where failure probability is negligible new equilibria never emerge, and that the social cost may decrease but it never increases. For the case of non-negligible failure probabilities, we provide a full characterization of the maximal impact of failures on the social cost under worst-case equilibrium outcomes.

Introduction

Congestion games (Rosenthal 1973) are a well-studied model of strategic sharing of resource, and have been used to investigate domains ranging from network design and routing (Kunniyur and Srikant 2003; Anshelevich et al. 2004) to cloud-computing and load-balancing (Suri, Tóth, and Zhou 2007; Vöcking 2007; Ashlagi, Tennenholtz, and Zohar 2010).

The characterization and computation of equilibrium outcomes in congestion games have received much attention (see e.g. (Fabrikant, Papadimitriou, and Talwar 2004; Jeong et al. 2005; Hayrapetyan, Tardos, and Wexler 2006; Ashlagi, Monderer, and Tennenholtz 2007)). In particular, researchers focused on the *Price of Anarchy*, which is the gap between the optimal cost and the cost under equilibrium outcome (Roughgarden and Tardos 2004; Christodoulou and Koutsoupias 2005). Nevertheless, an implicit assumption underlying all of this vast literature, is that agents who decided to use a certain resource always succeed in doing so. In practice, however, agents may fail to follow their chosen strategies, thereby utterly changing the costs of the game.

Consider a simple motivating example, where two travelers (our agents) wish to go from the airport to the city. Taxis to the city depart from gate C or gate E, where the taxis in gate E cost almost twice as much as those in gate C. The

travelers cannot communicate but if they happen to ride together, they share the cost of the ride equally. This can be modeled as a congestion game with two strategies, where (C, C) (sharing a cheap taxi) is optimal. However (E, E) is also an equilibrium. Consider what happens if both travelers know that their peer has some probability of failing to arrive, leaving the other to face the full costs of the ride (no matter what gate they may choose). In this new perturbed game it is a dominant strategy to take taxi from gate C (and hope that the other traveler will not fail to arrive, and choose the same gate). The “bad” equilibrium (E, E) dissolves.

Indeed, in most everyday interactions we cannot assume players are completely reliable. This is particularly true in computerized and online environments, where agents may inadvertently disconnect, face communication delays, etc. The above example shows that the equilibrium outcomes can change considerably when agents may fail, and that lack of reliability may lead to a more socially desirable outcome. These observations highlight the importance of understanding how failures affect the predicted outcomes of games.

We suggest a natural extension to the standard model of congestion games, which attributes a *survival probability* to each agent. Since in every congestion game the costs of players are determined only by the *number* of agents using a resource, it is straightforward to derive the new costs. In the absence of some agents, we compute the cost induced by the surviving agents, where each agent now aims to minimize her *expected cost* over all the realizations of the game.

Related work

Uncertainty in congestion games Though to the best of our knowledge no previous work studies the effects of agent failures on equilibria in congestion games, several works do examine similar themes. Penn et al. (2009; 2011) study congestion games with failure of *resources* rather than agents. In their model uncertainty always has a hazardous effect, as it encourages the agents to overload the system. While our model relies on the fact that congestion games already naturally define the costs for any set of surviving agents, Penn et al. must make specific assumptions regarding costs incurred when a resource fails.

A different model of uncertainty was introduced by Balcan et al. (2009), where agents perceive a noisy signal of the cost, which is either random or adversarial. Agents are un-

aware of the actual cost distribution, and are assumed to follow a myopic best-response strategy, which may lead them far away from any equilibrium. Balcan et al. study the *Price of Uncertainty* (PoU) in congestion games, which is the increase in social cost due to these perturbed dynamics.

Agent failures in games In general normal-form games there is no clear interpretation for a failure of an agent. However, there are particular families of games where failures do have a straightforward meaning. Messner and Polborn (2002) study how failures of voters to cast their vote shape the equilibria of election systems, focusing on the limit case where failure probability is negligible.

Closest in spirit to this paper is the work of Bachrach et al. (2011) which considers agent failures on *cooperative games* with transferable utilities. They prove that as in our case, failures in such games tend to have a beneficial effect. This is since failures can expand the core of the original game, thereby increasing its stability against collusion.

Our contribution

Our primary conceptual contribution is the introduction of agent failures to congestion games.

We first prove that every congestion game with failures always admits at least one pure Nash equilibrium, even if the induced game is *not* a congestion game. We then focus on a simpler scenario where each agent survives with a uniform independent probability p . We analyze both the limit behavior, where the survival probability goes to 1, and the case of fixed survival probabilities. In the limit case, we show that failures are beneficial: while the costs never increase, certain “bad equilibria” may be eliminated, thereby decreasing the worst social cost by an unbounded factor. Interestingly, we show that this no longer holds for Resource Selection games with increasing costs. For the case of fixed probabilities, we provide a full characterization of the maximal effect that failures may have on the Price of Anarchy, in terms of the probability p and the number of agents n . All omitted proofs can be found in the full version of this paper.¹

Definitions and Preliminaries

A **Congestion game** G is defined by a set of n agents N , and a set of resources F , each coupled with a cost function $c_j : [n] \rightarrow \mathbb{R}_+$. We denote the costs of resource $x \in F$ by a cost vector $c_x = (c_x(1), c_x(2), \dots, c_x(n))$. The highest possible cost on any single resource in a given game G is denoted by $M_G = \max_{x \in F, k \leq n} c_x(k)$. Each agent has a set of allowed strategies $S_i \subseteq 2^F$. A *strategy profile* is a vector of strategies $\mathbf{A} = (A_1, \dots, A_n)$, where $A_i \in S_i$. For every profile \mathbf{A} , each agent i incurs a cost (negative utility) $cost_i(G, \mathbf{A}) = \sum_{x \in A_i} c_x(n_x)$, where n_x is the number of agents that selected resource x in \mathbf{A} (including i). The *social cost* (or total cost) of a profile \mathbf{A} is: $cost(G, \mathbf{A}) = \sum_{i=1}^n cost_i(G, \mathbf{A}) = \sum_{x \in F} n_x c_x(n_x)$. We denote by $OPT(G)$ the minimal total cost over all profiles, i.e. $OPT(G) = \min_{\mathbf{A} \in \times_{i=1}^n S_i} cost(G, \mathbf{A})$. For simplicity,

¹Available from <http://tinyurl.com/bm5okau>.




Figure 1: The network K . An allowed strategy is a path from s to t , e.g. $A_1 = (s, x, t)$.

we assume all costs are non-negative integers, and (unless explicitly stated otherwise) that all costs are non-zero.²

Nash equilibrium A profile \mathbf{A} in G is a (pure) *Nash equilibrium* (NE) if no agent can gain by departing from \mathbf{A} : for any strategy $A'_i \in S_i$, $cost_i(G, \mathbf{A}) \leq cost_i(G, (A'_i, A_{-i}))$, where $-i = N \setminus \{i\}$. All congestion games are potential games, and thus admit a pure Nash equilibrium (Rosenthal 1973). In this work we restrict our attention to pure Nash equilibria.

Types of congestion games We focus on games where cost functions are either (weakly) *decreasing* or *increasing*. We denote such games by \check{G} or \hat{G} , respectively. Congestion games where all S_i are equal are called *symmetric*.

In a *resource selection game* (RSG), each agent i selects exactly one resource j from F . In *restricted resource selection games* (RRSG), which are an extension of RSGs, each agent i is restricted to select a single resource from $S_i \subseteq F$.

A different extension is *symmetric routing games* (SRTG) on a graph (V, E) , where each agent $i \in N$ selects a path from a source $s \in V$ to a target $t \in V$. An example of an SRTG (without the costs) is in Figure 1.

Note that in symmetric games (such as RSGs and SRTGs) with decreasing costs there is always an optimal NE where all agents select the same strategy.

Price of Anarchy The *Price of Anarchy* (PoA) of a game G compares the social cost of the worst Nash equilibrium to the optimal social cost, that is, $PoA(G) = \frac{cost(G, \mathbf{A}^*)}{OPT(G)}$, where \mathbf{A}^* is the *pure* NE with maximal cost in G .

Agent failures

Given a game G , we extend it with survival probabilities to every agent. In general, failures may be correlated, so we have a vector $\mathbf{p} \in \Delta(2^N)$, s.t. $p(S)$ is the probability that exactly the set S of agents survives to play. For any subsets $T \subseteq R \subseteq N$, let $p(T : R) = \sum_{S \subseteq N \setminus R} p(T \cup S)$, i.e. it denotes the probability that from all agents in R , exactly the agents of T survive. For any game G and a survival vector \mathbf{p} , we define the *reliability extension* $G^{\mathbf{p}}$ of G , by computing the *expected cost* that each surviving agent experiences.

If an agent j selects resource x , she is only affected by the failures of *other* agents on x . Thus, agent j will pay

$$c_{j,x}^{\mathbf{p}}(N_x) = \sum_{R \subseteq N_x \setminus \{j\}} p(R \cup \{j\} : N_x | j) c_x(|R| + 1),$$

²This technical assumption is required to avoid issues of division by zero when computing a ratio between costs.

where N_x is the set of agents selecting resource x . Note that if c_x is decreasing, then $c_{j,x}^{\mathbf{p}}(N_x) \geq c_x(N_x)$, and if c_x is strictly decreasing, then for all $|N_x| > 1$ the inequality is strict. Similarly, if c_x is increasing then $c_{j,x}^{\mathbf{p}}(N_x) \leq c_x(N_x)$.

In the general case the game $G^{\mathbf{p}}$ is *not* a congestion game, as the cost for agent j depends both on the identity of j , and on the identity of the other agents sharing the resource. One may wonder if this new game still has a pure Nash equilibrium, since this is *not* guaranteed in other extensions such as weighted congestion games (Milchtaich 1996). Our model may initially seem as an even broader generalization, as it allows dependencies among the agents. Somewhat surprisingly, we show that any reliability extension of G does admit a pure NE (see Theorem 1).

Games with i.i.d. failures In many cases we can avoid considering complicated failure distributions, and instead assume that each agent survives independently with a known probability $p \in (0, 1)$. In this case the number of surviving agents on each resource is simply a Binomial random variable. In particular, the cost to agent j does not depend on the identity of j . That is, all (surviving) agents on resource x pay $\mathbb{E}_{Z \sim \text{Bin}(n_x-1, p)}[c_x(Z+1)]$. Equivalently, $c_x^{\mathbf{p}}(n_x) = \sum_{k=0}^{n_x-1} \binom{n_x-1}{k} p^k (1-p)^{n_x-k} c_x(k+1)$.

We focus on measuring the effect that failures have on the game's outcome. For this purpose it is convenient to focus on i.i.d failures for two reasons: (a) they can be described by a single parameter p ; and (b) in contrast to the general case, the reliability extension $G^{\mathbf{p}}$ is also a congestion game.

Effect of failures on the costs Failures change costs in two distinct but interrelated ways.

Direct effect: the remaining players pay modified costs, as shown above. Note that the direct effect applies to optimal outcomes and to equilibrium outcomes alike. For example, if the costs in G are decreasing in the number of agents, then the direct effect of failures is that agents will now face higher costs in any given profile.

Indirect effect: the equilibria in the new game may change, leading to different payoffs.

We compute the total cost, summing over all the surviving players. That is,

$$\text{cost}(G^{\mathbf{p}}, \mathbf{A}) = \sum_{i=1}^n \text{cost}_i(G^{\mathbf{p}}, \mathbf{A}) = \sum_{x \in F} n_x c_x^{\mathbf{p}}(n_x).$$

We are particularly interested in cases where failure probabilities are low (i.e. when p is close to 1). In such cases the direct affect is negligible, but the indirect effect may play a major role. Specifically, we want to know if the equilibrium costs in the game can change significantly with small failure probabilities. When considering a ‘‘low probability’’ it is important to specify the order of quantifiers, i.e. whether the failure probability may depend on the game or not. In each result, we specify whether the survival probability p is allowed to take any *fixed* value. In contrast, when $p \rightarrow 1$ we can take an arbitrary value that may depend on the game. To demonstrate the difference, consider the following. For any fixed $p < 1$ there is a game $G = G(p)$ where $M_G > \frac{1}{1-p}$. However, for any fixed game G' , there is $p = p(G')$ sufficiently close to 1, s.t. $M_{G'} < \frac{1}{1-p}$.

We are mainly interested in the *indirect effect* of failures on the costs. To that end we compare the PoA of G and $G^{\mathbf{p}}$, which is a standard practice. Note that in the limit case $p \rightarrow 1$ the direct effect is negligible, so this is equivalent to measuring the indirect effect on the maximal costs.

General properties

We first prove that *any* reliability extension of a congestion game has a pure NE. We emphasize that no restriction on the cost functions is required for this result.

Theorem 1. *Let G be a congestion game, and \mathbf{p} a probability vector. Then $G^{\mathbf{p}}$ has a pure Nash equilibrium.*

Due to space constraints, we omit the full proof. However, it relies on the definition of the following function, which is a convex combination of the potential functions of all 2^n subgames of G .

$$\phi(\mathbf{A}) = \phi(N_1, \dots, N_{|F|}) = \sum_{R \subseteq N} p(R) \sum_{x \in F} \sum_{k=1}^{|R \cap N_x|} c_x(k).$$

While ϕ is *not* a potential function of $G^{\mathbf{p}}$, we show that it is *weighted potential function* of the game, where the weight of each agent is her own survival probability. Due to the existence of a weighted potential function, it is guaranteed that *any* sequence of best-replies by agents eventually converges to a pure Nash equilibrium (Monderer and Shapley 1996).

Another important issue is whether properties of the original game are conserved in $G^{\mathbf{p}}$. One property of interest is convexity (or concavity) of the cost functions, since such constraints can often be assumed in practice, and may have implications on the PoA. It turns out that convexity is maintained in the perturbed game (the proof is straightforward).

Proposition 2. *Let c_x be a convex [respectively, concave] cost function in the game G , and \mathbf{p} a probability vector. Then $c_{j,x}^{\mathbf{p}}$ is also convex [resp., concave], for all $j \in N$.*

In the remainder of this paper we assume that failures are i.i.d., that is that every player survives with probability p . We do note however, that most of our results easily extend to the more general cases of distinct (or correlated) probabilities.

Negligible failure probabilities

We now study how equilibria of a given game G are affected in the limit case. Most of the results assume negligible failure probabilities, but some hold for any $p < 1$ (e.g. Prop. 4).

Effect of failures on the set of NE

A crucial observation is that when failure probabilities are sufficiently low, no new NEs emerge caused by agent failure.

Proposition 3. *Let G be a congestion game. There is some $p^* = p^*(G)$ s.t. for all $p > p^*$, every NE profile of $G^{\mathbf{p}}$ is also an NE of G .*

Proof. If failure probabilities are negligible, then the costs in $G^{\mathbf{p}}$ can be arbitrarily close to the costs in G . Therefore all strict orders between costs remain, i.e. if $c_x(k) > c_y(k)$ then $c_x^{\mathbf{p}}(k) > c_y^{\mathbf{p}}(k)$. If $c_x(k) = c_y(k)$ then this equality

might break in G^p , but new equalities may not form. Finally, equality means that there is no incentive to deviate (from one strategy to another). Since equalities can only disappear, incentives to deviate can only increase, and Nash equilibria can only dissolve. \square

In contrast, the following examples demonstrate that certain NEs may dissolve even with a negligible failure probability, whether the costs are decreasing or increasing.

Proposition 4. *There is a RSG with decreasing costs \check{G}_1 and an NE \mathbf{A} in \check{G}_1 , s.t. for any survival probability $p < 1$, \mathbf{A} is not an NE of \check{G}_1^p .*

\check{G}_1 is an RSG with $n = 2$, $|F| = 2$, and we define costs as follows. $c_a = (M, 1)$ and $c_b = (M + 1, M)$, where $M > 1$. We can construct a similar example with increasing costs, by setting $c_a = (1, M)$, and $c_b = (M, 2M)$. Thus:

Proposition 5. *There is a RSG with increasing costs \hat{G}_1 and an NE \mathbf{A} in \hat{G}_1 , s.t. for any survival probability $p < 1$, \mathbf{A} is not an NE of \hat{G}_1^p .*

Effect of failures on the PoA

We show that if failure probabilities are small, the PoA cannot significantly increase.

Proposition 6. *Let G be a given congestion game with bounded PoA. For any $\varepsilon > 0$ there is $p^* = p^*(G, \varepsilon)$ s.t. for all $p \geq p^*$, $PoA(G^p) \leq PoA(G)(1 + \varepsilon)$.*

Proof sketch. We can set p^* arbitrarily close to 1. Therefore, by Prop. 3, there are no new equilibria in G^p . In particular, there are no new *bad* equilibria. Moreover, since costs are bounded, for every profile \mathbf{A} and agent i , $|cost_i(G^p, \mathbf{A}) - cost_i(G, \mathbf{A})|$ can be made arbitrarily small. Thus there is no indirect effect, and the direct effect is negligible for sufficiently small failure probabilities. \square

By Proposition 6 the PoA cannot increase due to failures. However the PoA might *decrease* due to the elimination of “bad” equilibria, and we would like to quantify this effect.

Decreasing costs In the RSG \check{G}_1 above one of the two NEs of the game dissolved when we added (even negligible) failure probabilities. Moreover, the removed NE was the worst NE in terms of social welfare. To be precise, without failures we had that $PoA(\check{G}_1) = M/1 = M$, whereas with failures the unique remaining NE was optimal, i.e. $PoA(\check{G}_1^p) = 1$. We get the following as a corollary,

Proposition 7. *For any M , there is a RSG with decreasing costs and two players \check{G}_1 s.t. (a) $PoA(\check{G}_1) > M$ (i.e. it is unbounded); and (b) for any $p < 1$, $PoA(\check{G}_1^p) = 1$.*

Increasing costs We next study the improvement in the PoA due to failures in games with *increasing costs*. The main result of this section is that in RSGs, i.e. symmetric singleton games, such a decrease is impossible. We first show that both symmetry and the singleton restriction are minimal. That is, if either one is relaxed, then there is an example where the PoA can improve arbitrarily.

Proposition 8. *For any M , there is a RSG with increasing costs and three players \hat{G}_2 s.t. (a) $PoA(\hat{G}_2) > M$; and (b) for any $p < 1$ $PoA(\hat{G}_2^p) = 1$.*

Proposition 9. *For any M there is an SRTG with increasing costs and two agents \hat{G}_3 (over the network K from Fig. 1), such that (a) $PoA(\hat{G}_3) = \Omega(M)$; and (b) for any $p < 1$, $PoA(\hat{G}_3^p) = 1$.*

RSGs with increasing costs To conclude this section, we show that when costs are increasing, the PoA can neither increase nor decrease due to negligible failure probabilities – in contrast to games with decreasing costs.

Lemma 10. *Let \hat{G} be a RSG with increasing costs. Let $c^* = cost(\hat{G}, \mathbf{A}^*)$ be the cost of the worst NE in \hat{G} . For any $p < 1$ there is another profile \mathbf{B} which is a pure NE in \hat{G}^p , and*

$$cost(\hat{G}, \mathbf{B}) \geq c^* - R_{\hat{G}} \cdot (1 - p),$$

where $R_{\hat{G}}$ is a constant that depends only on \hat{G} .

Proof. If \mathbf{A}^* is an NE in \hat{G}^p then we are done. Therefore assume that it is not, and thus there is an agent $i \in N$ which gains (in \hat{G}^p) by moving from some resource $a \in F$ to another $b \in F$. If there is more than one such deviation, then b is the strategy (resource) where i pays the lowest cost (break ties arbitrarily). Denote by \mathbf{A}_1 the outcome where i plays b instead of a , and all other agents play as in $\mathbf{A}_0 \equiv \mathbf{A}^*$. As long as \mathbf{A}_t is not an NE (in \hat{G}^p), we repeat the process until no agent wants to deviate, and denote the final profile by \mathbf{B} . We argue that there are at most n steps until convergence.

If an agent i moves from a to b in step t then no agent will leave resource b in a future step $t' > t$ (otherwise agent i would have had a better step at time t). Thus there are mutually exclusive subsets $A, B \subseteq F$ s.t. agents only move from A to B . In particular, this means that each agent moves at most once and thus there are at most n steps.

Let $M = M_{\hat{G}}$ (a constant). We next show that for all t , $\delta_t \equiv cost(\hat{G}, \mathbf{A}_{t-1}) - cost(\hat{G}, \mathbf{A}_t) \leq O(n^2 M(1 - p))$.

We denote by n_j^*, n_j, n'_j the number of agents using resource j in the profiles $\mathbf{A}^*, \mathbf{A}_{t-1}$ and \mathbf{A}_t , respectively. Suppose that between \mathbf{A}_{t-1} and \mathbf{A}_t some agent i moved from a to b . Then $n_a^* \geq n_a = n'_a + 1$ and $n_b^* \leq n_b = n'_b - 1$. Since \mathbf{A}^* is an NE in \hat{G} , and by monotonicity of c_j ,

$$c_a(n_a) \leq c_a(n_a^*) \leq c_b(n_b^* + 1) \leq c_b(n_b + 1) = c_b(n'_b). \quad (1)$$

On the other hand, since i preferred b over a in \hat{G}^p ,

$$c_a^p(n_a) > c_b^p(n'_b). \quad (2)$$

We next bound the two expressions. Denote $\alpha = 1 - p$. Denote $\Delta_a = c_a(n_a) - c_a(n_a - 1)$, and $\Delta_b = c_b(n'_b) - c_b(n'_b - 1)$. There is a probability of $p^{n_a - 1} < 1 - (n_a - 1)\alpha + (n_a - 1)^2\alpha^2$ that all agents on a (except i) survive. Thus w.p. of at least $(n_a - 1)\alpha - (n_a - 1)^2\alpha^2$ at least one agent fails. Thus

$$c_a^p(n_a) \leq c_a(n_a) - ((n_a - 1)\alpha - (n_a - 1)^2\alpha^2)\Delta_a. \quad (3)$$

Similarly, the probability that exactly one agent fails in resource b is at most $(n'_b - 1)\alpha = n_b\alpha$ (in which case the cost

drops by Δ_b), and the probability that more than one agent fails is at most $n_b^2 \alpha^2$ (in which case the cost drops by at most M). Thus $c_b^p(n_b') \geq c_b(n_b') - n_b \alpha \Delta_b - (n_b \alpha)^2 M$.

By combining the last equation with Eq. (1),(2) and (3) ,

$$\begin{aligned} c_a(n_a) - ((n_a - 1)\alpha - (n_a - 1)^2 \alpha^2) \Delta_a &\geq \\ c_b(n_b') - n_b \alpha \Delta_b - (n_b \alpha)^2 M &\geq c_a(n_a) - n_b \alpha \Delta_b - (n_b \alpha)^2 M \end{aligned}$$

Then, by rearranging terms,

$$\begin{aligned} n_b \Delta_b + (n_b)^2 \alpha M &\geq ((n_a - 1) - (n_a - 1)^2 \alpha) \Delta_a \quad \Rightarrow \\ n_b \Delta_b &\geq (n_a - 1) \Delta_a - (n_a - 1)^2 \alpha \Delta_a - (n_b)^2 \alpha M \\ &\geq (n_a - 1) \Delta_a - 2\alpha n^2 M \end{aligned} \quad (4)$$

We can now bound the costs of $\mathbf{A}_{t-1}, \mathbf{A}_t$.

$$\begin{aligned} \delta_c &= n_a c_a(n_a) + n_b c_b(n_b) - (n_a' c_a(n_a') + n_b' c_b(n_b')) \\ &= n_a' (c_a(n_a) - c_a(n_a')) + c_a(n_a) - n_b (c_b(n_b') - c_b(n_b)) - c_b(n_b') \\ &= (n_a - 1) \Delta_a - n_b \Delta_b + (c_a(n_a) - c_b(n_b')) \\ &\leq (n_a - 1) \Delta_a - n_b \Delta_b \leq 2\alpha n^2 M, \end{aligned} \quad (\text{by (1),(4)})$$

Finally, since there are at most n steps, we get that

$$\text{cost}(\widehat{G}, \mathbf{B}) \geq c^* - n \cdot (2\alpha n^2 M) = c^* - R_{\widehat{G}}(1 - p). \quad \square$$

Proposition 11. *Let \widehat{G} be a RSG with increasing costs. Then for any $\varepsilon > 0$ there is some $p < 1$ s.t. the ratio between $\text{PoA}(\widehat{G})$ and $\text{PoA}(\widehat{G}^p)$ is small, i.e.*

$$\text{PoA}(\widehat{G})(1 - \varepsilon) \leq \text{PoA}(\widehat{G}^p) \leq \text{PoA}(\widehat{G})(1 + \varepsilon).$$

Proof sketch. The crux of the proof is Lemma 10, showing that although *some* bad equilibria may dissolve in G^p , at least one bad equilibrium (that is $\varepsilon/3$ close to the worst equilibrium \mathbf{A}^*) survives if p exceeds some value p^* .

We then set p high enough so that (a) For every profile \mathbf{A} , $\text{cost}(G^p, \mathbf{A}) \geq \text{cost}(G, \mathbf{A}) - \varepsilon/3$ (i.e. the direct effect is negligible); (b) No new equilibria emerge (i.e. Prop. 6 holds); and (c) $p > p^*$ (i.e. $(1 - p)R_{\widehat{G}} < \varepsilon/3$).

Since $\text{OPT}(\widehat{G}) > 0$, then it is at least 1 as all costs are integers. Then by (c) and Lemma 10 there is a bad equilibrium \mathbf{B} that still exists in G^p , and by (a) both OPT and the cost of \mathbf{B} do not improve much in \widehat{G}^p . Thus

$$\begin{aligned} \text{PoA}(\widehat{G}^p) &= \frac{\text{cost}(\widehat{G}^p, \mathbf{B}^*)}{\text{OPT}(\widehat{G}^p)} \geq \frac{\text{cost}(\widehat{G}, \mathbf{B})(1 - \varepsilon/3)}{\text{OPT}(\widehat{G})(1 + \varepsilon/3)} \\ &\geq \frac{\text{cost}(\widehat{G}, \mathbf{A}^*)(1 - \varepsilon/3)(1 - \varepsilon/3)}{\text{OPT}(\widehat{G})(1 + \varepsilon/3)} \\ &= \text{PoA}(\widehat{G}) \frac{(1 - \varepsilon/3)^2}{1 + \varepsilon/3} \geq \text{PoA}(\widehat{G})(1 - \varepsilon). \end{aligned}$$

The upper bound follows directly from (b). \square

Fixed failure probabilities

In this section we assume that there is some fixed survival probability p , whereas the parameters of the game may vary. Interestingly, it turns out that fixing the probability before the game is defined (i.e. changing the order of quantifiers) is

highly significant, and some results are very different from the ones in the previous section. Recall for example that when $p \rightarrow 1$, it was impossible to introduce new NEs to a game via failures. However this is no longer true when p is fixed (even if small), and the costs may significantly vary.³

Effect of failure on the set of NEs

While some NEs may disappear, no new NEs can emerge in a symmetric game with decreasing costs.

Proposition 12. *Let \check{G} be a symmetric game with decreasing costs, and let $p < 1$. Then \check{G}^p does not admit new Nash equilibria.*

However, symmetry turns out to be a minimal requirement. Note that the game \check{G}_2 depends on the value of p .

Proposition 13. *For any $p < 1$ there is a RMSG with two agents and decreasing costs \check{G}_2 s.t. \check{G}_2^p has new NEs.*

As for games with *increasing costs*, they can behave quite differently from games with decreasing costs when there are fixed failure probabilities (even small ones). In particular, new NEs may emerge even in symmetric games.

Proposition 14. *For any $p < 1$, there is a RSG with increasing costs \widehat{G}_4 , such that \widehat{G}_4^p has new NEs.*

Example. The game \widehat{G}_4 has two resources $\{a, b\}$ and n agents. a always costs $M > 1$. b costs 1, unless everybody select it, and then it costs $R > M$. \diamond

Effect on the PoA – Games with decreasing costs

It is quite clear that with significant failure probabilities, the social cost of playing some NE in a game may increase. However since the cost of OPT may also increase, it is not clear how the PoA is affected. The following examples show that PoA can increase as well – in contrast to the result we had when failure probabilities are negligible.

Proposition 15. *For any M and any $p < 1$, there is a RMSG \check{G}_2 with three players s.t. (a) $\text{PoA}(\check{G}_2) = 1$; and $\text{PoA}(\check{G}_2^p) > M$.*

That is, in asymmetric games we can get an unbounded increase in the PoA (in fact, \check{G}_2 is the same game used in Prop. 13). When \check{G} is symmetric, there is a tight bound on the PoA – and thus on the maximal increase in the PoA.

Proposition 16. *Let \check{G} be a symmetric game with decreasing costs. For any $p < 1$ it holds that $\text{PoA}(\check{G}^p) \leq (1 - p)^{1-n}$.*

Proposition 17. *For any $p < 1$, any n , and any $\varepsilon > 0$, there is a RSG with decreasing costs \check{G}_3 s.t. (a) $\text{PoA}(\check{G}_3) = 1$; and (b) $\text{PoA}(\check{G}_3^p) \geq (1 - p)^{1-n} - \varepsilon$.*

Example. The game \check{G}_3 contains n players and 2 resources with the following costs: $c_a = (M, 1, 1, \dots, 1)$, and $c_b = (R, \dots, R, R, 1)$, where $R = \frac{M - p^{n-1}}{1 - p^{n-1}}$. \diamond

The bound of $(1 - p)^{1-n}$ is somewhat counter-intuitive. For a fixed game \check{G} , we know that increasing the survival

³To see these contrasts more clearly, the reader is advised to look at Tables 1, 2 and 3 in the last section.

probability p eventually means that the PoA cannot increase (much). It therefore seems reasonable to assume that this effect is “monotone”, i.e. that as p grows, then the maximal ratio $\frac{PoA(\hat{G}^p)}{PoA(\hat{G})}$ becomes smaller and smaller. However, the converse is true: While for small p the ratio is also small, as p grows we can find examples where this ratio becomes larger and larger.

Another interesting implication is that the PoA of \hat{G}^p is bounded, whereas this is not true for games without failures. Some insight might be gained by the following explanation. The cost of the worst equilibrium can sharply increase for any probability. However, for low p a high increase must entail that the *optimal cost* is also increasing, thereby limiting the maximal ratio between the two.

Effect on the PoA – Games with increasing costs

Lemma 18. *For any RSG with increasing costs \hat{G} , $1 \leq PoA(\hat{G}) \leq n$.*

In particular, the lemma entails that the PoA of \hat{G} can never decrease or increase by a factor of more than n .

Bounds on the increase in PoA By properly setting the parameters of the game \hat{G}_4 (from Prop. 14), we get:

Proposition 19. *For any $p < 1$, any $\varepsilon > 0$ and any number of players n , there is a RSG with increasing costs \hat{G}_4 , s.t. (a) $PoA(\hat{G}_4) = 1$; and (b) $PoA(\hat{G}_4^p) > n - \varepsilon$.*

If we either relax the symmetry constraint, or allow more complex strategies than singletons, then the PoA may increase by an unbounded factor (examples omitted).

Proposition 20. *For any $\frac{1}{2} < p < 1$ and any constant M , there is a RSG with increasing costs and three players \hat{G}_5 s.t. (a) $PoA(\hat{G}_5) = 1$; and (b) $PoA(\hat{G}_5^p) > M$.*

Proposition 21. *For any $p < 1$ and M , there is a SRTG \hat{G}_6 with increasing costs and four players s.t. (a) $PoA(\hat{G}_6) = 1$; and (b) $PoA(\hat{G}_6^p) > M$.*

Example. Set R s.t. $R > 2M/p^3$ and $\frac{R}{R+7} > p$ (for $p > \frac{1}{2}$). Consider the SRTG network K from Figure 1, with the costs as follows. $c_{(x,y)} = (1, 1, 1, R+8)$, and the cost of the other four edges is $(1, 1, R+5, R+5)$. \diamond

Bounds on lowering the PoA Prop. 9 shows that failures can trigger an unbounded improvement in the PoA in routing games, even if they are symmetric. Our last result concludes that with fixed failure probabilities even the PoA of RSGs can improve, although not by an unbounded factor.

Proposition 22. *Suppose $1 > p > \frac{1}{2}$. There exists a family of RSG (with $n = 2, 3, 4, \dots$ agents) with increasing costs \hat{G}_7 , s.t. (a) $PoA(\hat{G}_7) = \Omega(n)$; and (b) $PoA(\hat{G}_7^p) = O(1)$.*

Discussion

Two particular conclusions can be drawn from our results. First, failures may completely alter the outcome of the game, even if they occur with a very low probability. Thus they

Decreasing costs	NE may dissolve	NE may emerge	
		symmetric	any game
$p < 1$	yes (\uparrow)	no (P. 12)	yes (P. 13)
$p \rightarrow 1$	yes (P. 4)	no (\downarrow, \Leftarrow)	no (P. 3)
increasing costs			
$p < 1$	yes (\uparrow)	yes (P. 14)	yes (\Rightarrow)
$p \rightarrow 1$	yes (P. 5)	no (\Leftarrow)	no (P. 3)

Table 1: The table describes how NEs in G^p may differ from those in G . “yes” means that there is an example where the described effect occurs. P. # refers to Proposition #.

Dec. costs	Max. decrease in PoA	Maximal increase in PoA	
		symmetric	other
$p < 1$	UB (P. 7)	$(1-p)^{1-n}$ (*)	UB (P. 15)
$p \rightarrow 1$	UB (P. 7)	none (\Leftarrow)	none (P. 6)

Table 2: The table describes the bounds on the maximal ratio between $PoA(\hat{G}^p)$ and $PoA(\hat{G})$. “none” means there is no change, or effect is negligible. “UB” means the change is unbounded in terms of p and n . (* by P. 17 and P. 16)

must be taken into account in the analysis of many realistic scenarios. Second, some limited level of noise (in the form of failures) can actually contribute to the participating players, by eliminating bad equilibria. Two notable examples for this are Prop. 7 showing an unbounded improvement in the social cost; and Prop. 16 showing an upper bound on the PoA of whole family of games, where no such bound exists for games without failures.

Concavity and convexity In many realistic games we can assume that marginal costs are increasing or decreasing. We have shown that this property does not change when failures occur. However concavity/convexity can potentially limit the PoA or the ratio by which the PoA changes due to failures. We note that all our results for the limit case hold regardless of convexity or concavity. However, some examples in the latter section make use of particular cost functions. We leave it as an open question whether convex/concave examples can be constructed in each case.

Future Work Many questions are left open for future research. These include understanding the effect of failures on the *best Nash equilibria* (e.g. by studying the Price of Stability); focusing on particular interesting families of cost functions; and bounding the rate of convergence of various game dynamics. We also believe that with strictly monotone cost functions (and in particular convex or concave families) some of our results may change.

An important future goal is to leverage our current knowledge on uncertainty in congestion games in various models, to prompt the design of better *mechanisms*. That is, to intelligently manipulate the reliability of the connections or the information players have on the number of survivors, so as to benefit the society by eliminating unwanted equilibria.

Inc. costs	Maximal increase in PoA		
	RSG	symmetric	other
$p < 1$	n (P. 19, L. 18)	UB (P. 21)	UB (P. 20)
$p \rightarrow 1$	none (\Leftarrow)	none (\Leftarrow)	none (P. 6)
Maximal decrease in PoA			
$p < 1$	$\Theta(n)$ (P. 22, L. 18)	UB (P. 9)	UB (\Rightarrow)
$p \rightarrow 1$	none (P. 11)	UB (P. 9)	UB (\Rightarrow)

Table 3: (see caption of Table 2).

References

- Anshelevich, E.; Dasgupta, A.; Kleinberg, J.; Tardos, E.; Wexler, T.; and Roughgarden, T. 2004. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, FOCS '04*, 295–304.
- Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2007. Learning equilibrium in resource selection games. In *Proceedings of the 21st National Conference on Artificial Intelligence, AAAI '07*, 18–23.
- Ashlagi, I.; Tennenholtz, M.; and Zohar, A. 2010. Competing schedulers. In *Proceedings of the 24th National Conference on Artificial Intelligence, AAAI '10*, 691–696.
- Bachrach, Y.; Meir, R.; Feldman, M.; and Tennenholtz, M. 2011. Solving cooperative reliability games. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence, UAI '11*, 27–34.
- Balcan, M.-F.; Blum, A.; and Mansour, Y. 2009. The price of uncertainty. In *Proceedings of the 10th ACM conference on Electronic commerce, EC '09*, 285–294.
- Christodoulou, G., and Koutsoupias, E. 2005. The price of anarchy of finite congestion games. In *Proceedings of the 37th annual ACM symposium on Theory of computing, STOC '05*, 67–73.
- Fabrikant, A.; Papadimitriou, C.; and Talwar, K. 2004. The complexity of pure nash equilibria. In *Proceedings of the 36th annual ACM symposium on Theory of computing, STOC '04*, 604–612.
- Hayrapetyan, A.; Tardos, É.; and Wexler, T. 2006. The effect of collusion in congestion games. In *Proceedings of the 38th annual ACM symposium on Theory of computing, STOC '06*, 89–98.
- Ieong, S.; McGrew, R.; Nudelman, E.; Shoham, Y.; and Sun, Q. 2005. Fast and compact: A simple class of congestion games. In *Proceedings of the 19th National Conference on Artificial Intelligence*, volume 20 of AAAI '05, 489.
- Kunniyur, S., and Srikant, R. 2003. End-to-end congestion control schemes: Utility functions, random losses and ecn marks. *Networking, IEEE/ACM Transactions on* 11(5):689–702.
- Messner, M., and Polborn, M. K. 2002. Robust political equilibria under plurality and runoff rule. Mimeo, Bocconi University.
- Milchtaich, I. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13(1):111–124.

- Monderer, D., and Shapley, L. S. 1996. Potential games. *Games and Economic Behavior* 14(1):124–143.
- Penn, M.; Polukarov, M.; and Tennenholtz, M. 2009. Congestion games with load-dependent failures: Identical resources. *Games and Economic Behavior* 67(1):156 – 173.
- Penn, M.; Polukarov, M.; and Tennenholtz, M. 2011. Congestion games with failures. *Discrete Applied Mathematics* 159(15):1508 – 1525.
- Rosenthal, R. 1973. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2:65–67.
- Roughgarden, T., and Tardos, E. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47(2):389 – 403.
- Suri, S.; Tóth, C.; and Zhou, Y. 2007. Selfish load balancing and atomic congestion games. *Algorithmica* 47(1):79–96.
- Vöcking, B. 2007. Selfish load balancing. In Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 20, 517–542.