

Risk Sensitivity of Price of Anarchy under Uncertainty

GEORGIOS PILIOURAS, Singapore University of Technology and Design

EVDOKIA NIKOLOVA, University of Texas at Austin

JEFF S. SHAMMA, Georgia Institute of Technology and King Abdullah University of Science and Technology

In game theory, the price of anarchy framework studies efficiency loss in decentralized environments. Optimization and decision theory, on the other hand, explore tradeoffs between optimality and robustness in the case of single agent decision making under uncertainty. What happens when we combine both approaches?

We examine connections between the efficiency loss due to decentralization and the efficiency loss due to uncertainty and establish tight performance guarantees for distributed systems in uncertain environments. We present applications based on novel variants of atomic congestion games with uncertain costs, for which we provide tight performance bounds under a wide range of risk attitudes. Our results establish that the individual's attitude towards uncertainty has a critical effect on system performance and therefore should be a subject of close and systematic investigation.

CCS Concepts: • **Theory of computation** → **Quality of equilibria**;

Additional Key Words and Phrases: Price of Anarchy, Stochastic Optimization, Congestion Games

ACM Reference Format:

Georgios Piliouras, Evdokia Nikolova, and Jeff S. Shamma, 2016. Risk Sensitivity of Price of Anarchy under Uncertainty *ACM Trans. Econ. Comput.* V, N, Article A (January YYYY), 29 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Ever since its introduction by Koutsoupias and Papadimitriou [1999], the price of anarchy framework has played a key role in the formalized study of efficiency loss in competitive environments due to the lack of centralized control. Defined as the ratio between the social welfare at the worst Nash equilibrium and its optimal value, the price of anarchy embodies several desirable attributes. Conceptually simple and interesting, concise in representation, and analytically tractable, as it has proven to be in a wide range of settings, the price of anarchy has had a significant impact on the way we perceive and argue about systems.

This work is supported by AFOSR/MURI Project FA9550-09-1-0538, AFOSR Project FA9550-09-1-0420, ONR Grant N00014-09-1-0751 and NSF CCF Grant 1216103, 1350823, 1331863, the CMI Wally Baer and Jeri Weiss postdoctoral fellowship, SUTD grant SRG ESD 2015 097 and MOE AcRF Tier 2 Grant “A dynamical systems approach to algorithmic game theory”. Part of the work was completed while Georgios Piliouras was a CMI Wally Baer and Jeri Weiss postdoctoral fellow at California Institute of Technology. Part of the work was completed while Georgios Piliouras and Evdokia Nikolova were resident scientists at the Simons Institute for the Theory of Computing. See the acknowledgements sections before the references.

Authors' addresses: G. Piliouras, Engineering Systems and Design Pillar, Singapore University of Technology and Design, georgios@sutd.edu.sg; E. Nikolova, Department of Electrical and Computer Engineering, University of Texas at Austin nikolova@austin.utexas.edu; J.S. Shamma, School of Electrical and Computer Engineering, Georgia Institute of Technology, shamma@gatech.edu and CEMSE Division, King Abdullah University of Science and Technology (KAUST), jeff.shamma@kaust.edu.sa.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© YYYY ACM. 1539-9087/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

Naturally, any theoretical analysis of real-world social or technological networks introduces simplifying assumptions that potentially could lead to significant divergence between theoretical predictions and actual system behavior. In fact, a key aspect of any such systematic investigation is distinguishing those modeling parameters that can safely be ignored from those that have a dominant impact and carefully exploring the latter subspace. Such concerns about modeling imprecisions are especially relevant when it comes to the study of socioeconomic phenomena since the “ground truth” of human nature and decision making is inherently complex and intractable in its full generality.

Nevertheless, during the last 70 years, a long line of work spanning economics, psychology, philosophy, mathematics, and statistics has been focused exactly on capturing the values, uncertainties and other primary factors relevant to rational decision making. Modern decision theory along with stochastic/robust optimization have codified different approaches to decision making under uncertainty. These models mark a wide range of departures from expectation maximization. Some of the models are inspired by the realities of everyday decision making and our limitations in perceiving accurately and responding optimally to all the subtleties of an uncertain environment. Other approaches are focused on the other extreme of the spectrum of cognitive abilities. Their goal is to design efficient algorithmic procedures which, based on computational preprocessing and intricate estimation techniques, allow for provably optimal nontrivial performance guarantees. For example, a company might wish to identify a strategic course of action that guarantees for itself the highest possible revenue for the coming year outside a 1% window of the worst case realizations of uncertainty, *i.e.*, catastrophic scenarios. This exploration of the tradeoffs between robustness and optimality, sometimes referred to as price of robustness [Aghassi and Bertsimas 2006; Bertsimas et al. 2011], has been the subject of intensive research within modern optimization theory.

These decision making models are not merely relevant to game theory but they were introduced as a direct response to the theoretical challenge of coupled decision making. In any game theoretic setting, even in the classic full information case, given any mixed strategy profile, the agents are required to solve an optimization problem under uncertainty. The seminal works by von Neumann and Morgenstern [1944], Nash [1951] and Harsanyi [1967] have laid the expectation maximization framework on such strong foundations within the game theory community that only recently have early steps been taken away from it [Aghassi and Bertsimas 2006; McKelvey and Palfrey 1995]. Even these works, however, are primarily definitional, focusing on establishing fixed point theorems under different decision theoretic models. In contrast, the price of anarchy framework seems like an ideal testing ground for a quantitative study of the implications of merging decision theory under uncertainty and game theory. In this work, we wish to show how settings with uncertain payoffs coupled with risk-aware agents may indeed not only be analytically tractable but furthermore hold the promise of revealing novel and sometimes surprising insights about well-studied competitive settings.

Our work focuses on a variant of a classic class of games within the price of anarchy framework, atomic congestion games with affine cost functions [Christodoulou and Koutsoupias 2005]. Such games abstractly capture the effects of everyday routing decisions and how they get translated to congestion on the road. In order to intuitively understand how different decision theoretic protocols are a natural and integral part of this setting, let’s consider an actual commuter deciding upon a route from her home to her office.

On any given morning, the driver, Alice, might be indifferent between arriving at her office a few minutes earlier or later, so picking the route with the minimal expected la-

Table I. Price of anarchy of linear congestion games under uniform schedulers.

Principle	Price of Anarchy
Minimizing Expected cost	$5/3$
Average case analysis	$5/3$
Win-or-Go-Home	unbounded
Second moment method	unbounded
Wald's minimax cost	2
Savage's minimax regret	$[4/3, 3]^1$

tency seems like a reasonable choice. Of course, even estimating the actual expected latency is a task that requires some mental overhead of remembering unexpected events in the past that affected her arrival time (*e.g.*, heavy rain with .01%, traffic accidents .02%). An easier alternative is to implicitly assume that this is a “typical day”. This average case analysis smooths out atypical events and allows for effortless decision making. On other occasions (*e.g.*, when giving a presentation at work), the same individual, faced with the same optimization problem, is now much more sensitive to unfortunate rare events. In this case, Alice may subconsciously overestimate the possibility of bad events, heavily penalizing routes with high variance of arrival times. In fact, in the case where on time arrival is absolutely critical, she might be willing to go through a lengthy detour of backroads where congestion effects are not a factor. Finally, Alice's everyday challenge is furthermore complicated by her coworker and neighbor, Bob. Bob is rather vocal about his uncanny prowess in picking fast routes. These differences between her actual arrival time and that of the best route in hindsight are sometimes the cause of regret and may affect future decisions. Of course, these effects are only aggravated on rainy days, since as fate would have it, there is only a single parking spot outside their office and whoever gets there first takes it whereas the other one must walk in the rain. A race it is then.

The paragraph above hopefully conveys that we can expect a wide range of attitudes towards uncertainty to materialize in our everyday experiences. Moreover, these attitudes are not necessarily tied to the nature of the specific optimization problem. They are not even necessarily hard-coded personality traits of different individuals, but may depend on the context that is not even present in the mathematical formulation of the setting. What are the effects of this interwoven net of unspoken and sometimes unconscious attitudes to the hard bottom-line of the realized system performance? Interestingly, and perhaps distressingly, our results formally show that this single risk parameter can completely dictate system performance. Attitude is, indeed, everything.

We introduce uncertainty in congestion games by inserting a randomizing scheduler on each congested element. The latency experienced by any user on a machine is a random variable depending both on the load of the machine and on the random ordering of the individual. Intuitively, similarly to the case of a printer that receives a number of requests simultaneously or to the case of a number of cars that merge from two lanes into one, the jobs/users get processed in some order and each job gets burdened only by those requests that get assigned ahead of it in the queue. In our default model, we will assume that the random order is unique across all resources so that if user i is processed before user j on the first edge of a common path, this ordering will persist

¹These bounds hold in the case of symmetric load balancing games. We also show a stronger lower bound of 2 in the general case.

throughout the path.² A nice feature of this approach is that, in any outcome of the game, the total latency experienced on any element (and on the system as a whole) is deterministic and in fact equal to the standard definition of potential in congestion games. This is indeed the case as on each element there will be exactly one user served first, one served second, and so on.³ Therefore, *regardless of the risk attitudes of the agents*, the optimal social cost is uniquely defined. Changes in the price of anarchy can only reflect the fact that different risk attitudes give rise to different stable states. As Table I indicates, our analysis provides tight bounds on the price of anarchy in the case of linear latency functions under a wide range of decision making principles.⁴

In the rest of this paper, we provide an overview of work on decision making under uncertainty as well as earlier price of anarchy related work. Section 3 contains a detailed exposition of the single agent decision making framework, whereas Section 4 combines uncertain decision making with games of uncertain payoffs, explicitly focusing on the case of congestion games with randomized scheduling policies. Sections 5-10 present the price of anarchy analysis for each of the following decision making principles: risk-neutral stochastic optimization (*i.e.*, optimizing expectation), average case analysis, win-or-go-home principle, second moment stochastic optimization, Wald's minimax principle and Savage's minimax regret principle for linear congestion games. Section 11 concludes with discussion about future directions on the intersection of decision theory under uncertainty and game theory. Finally, in Appendix D we present upper bounds on the efficiency of these models under polynomial latency functions. Unlike our analysis in case of linear latency functions these bounds are not tight. Nevertheless, these results already suffice to establish that the sensitivity of price of anarchy to uncertainty is not an artifact of the linearity of the latency functions but is indeed a more general phenomenon. Establishing tight price of anarchy bounds for all risk attitudes under polynomial cost functions is an interesting open question for future work.

2. RELATED WORK

There are many different ways to make decisions in uncertain environments, and consequently, to model a game in which the agents need to make such decisions. Stochastic programming is an entire field (*e.g.*, see textbooks [Prékopa 1995; Birge and Louveaux 1997; Shapiro et al. 2009]), yet it is just one approach to optimization under uncertainty. It is most commonly used to capture risk-neutral agents as its typical objective function is expected cost. Robust optimization (*e.g.*, [Ben-Tal et al. 2009; Bertsimas et al. 2011]), at the other end of the spectrum, is suitable for infinitely risk-averse agents as it seeks the best possible worst-case outcome. The robust framework has been the focus of a lot of work within operations research and related fields.

Risk-averse optimization aims to bridge the gap between the two extreme risk attitudes, by considering more moderate and realistic models of risk. In fact, this is not yet a well-defined field, since multiple competing approaches to modeling risk are pursued in different communities. In economics, a common approach for modeling risk is via utility functions [Neumann and Morgenstern 1944] and, consequently, various

²Our results carry over to several variants of this model. For a full description of the different models of randomizing schedulers see section 4.2.

³A complete discussion of the social cost definition and justification can be found in section 4.4.

⁴We should note here that our results do not critically depend on the assumption that the ordering of each job on each machine is chosen uniformly at random. For example, our results under Wald's minimax cost principle carry over even in cases of globally correlated (possibly adversarial) scheduling policies. We choose not to focus on these technical facts in the introduction in favor of a more uniform and simple exposition. See section 4.2 for more details.

alternatives to expected utility theory have been proposed, including weighted utility theory, prospect theory and the theory of ambiguity aversion [Machina 1987]. In finance, the mathematical study of risk for portfolio selection started with the mean-variance framework of Markowitz [1952]. More recently, due to a variety of paradoxes and pitfalls of the traditional approaches, an axiomatic approach to risk has been proposed [Artzner et al. 1999; Rockafellar 2007], considering coherent and convex risk measures. Such risk models have yet to be incorporated in a systematic way into formal algorithmic study and algorithmic game theory models.

A noteworthy special case of risk preference is captured in chance-constrained optimization (e.g., [Uryasev 2000; Boş et al. 2007]), where constraints with uncertain variables are required to hold with at least a given probability. Algorithmic work on risk for combinatorial problems has considered a variety of models, including mean-risk optimization, value-at-risk and probabilistic (chance) constraints (e.g., [Nikolova et al. 2006; Srinivasan 2007; Nikolova 2010; Swamy 2011]).

Games of incomplete information have been the subject of recent focus within the price of anarchy literature [Caragiannis et al. 2012; Paes Leme and Tardos 2010; Christodoulou et al. 2008; Lucier and Paes Leme 2011; Lucier and Borodin 2010; Syrgkanis 2012; Gairing and Tiemann. 2005; Garg and Narahari 2005; Georgiou et al. 2006]. Typically, the randomness is introduced via sampling the agent types randomly and independently from a commonly known distribution. Furthermore, the assumption of risk neutral agents is ubiquitous. Roughgarden has shown that the analysis of such settings essentially reduces to the deterministic case [Roughgarden 2012]. In our work, types are dependent and furthermore our results in the case of robust optimizing agents trivially extend even to cases where the agents hold different beliefs about the underlying type distribution. Moreover, our focus is on breaking away from the convention of risk-neutral agents and exploring the implications of different risk attitudes.

The idea of introducing randomized scheduling policies on each congested element was inspired by machine scheduling ([Heydenreich et al. 2007], and references therein). However, the two settings are quite distinct. Typically, in machine scheduling the set of strategies for each job consists of singletons (*i.e.*, individual machines), whereas our setting has a nontrivial topology (routing network). On the other hand, in machine scheduling jobs come in different sizes and the completion time on each machine can in the most general case be an arbitrary function of each (job, machine) pair. More importantly, the end goals of the two approaches are decidedly different. Mechanism design for machine scheduling aims at designing simple scheduling policies such that risk neutral agents converge to almost optimal outcomes. In our setting, the schedulers are used as an abstraction to model uncertainty. They are used to infuse uncertainty, sometimes even in an adversarial and correlated manner. The goal is not to guide the users but to provide a challenging setting so as to test the efficiency of different risk-sensitive decision models.

Little attention has been paid so far within algorithmic game theory to cases where the agents are not expectation maximizers. Fiat and Papadimitriou [2010] consider games with deterministic payoffs and incorporate risk resulting from mixed strategies via mean-variance and other objectives. Mavronicolas and Monien [2012] also consider mean-variance objectives for mixed equilibria and prove a PLS-completeness result for computing such equilibria. Perhaps most closely related to our model of uncertain payoffs is the stochastic selfish routing model of Nikolova and Stier-Moses where the edge latency functions are independent random variables with finite first and second moments [Nikolova and Stier-Moses 2011, 2014b,a].

Aghassi and Bertsimas [2006] present a robust game theory model where the payoffs come from bounded uncertainty sets and the agents assume adversarial realizations

of uncertainty. They show that fixed point theorems imply existence results for their equilibrium notion. Our model of win-lose games was inspired by Immorlica et al. [2011]. At a high level, the interplay of the decision making process of a single agent and the price of anarchy is related to the work on lookahead search in game playing by Mirrokni et al. [2012]. Back to uncertain payoffs, stochastic games from game theory, first introduced by Shapley [1953], refer to dynamic games in which agents make probabilistic transitions (*e.g.*, see [Neyman and Sorin 2003; Filar and Vrieze 1996]). They are a generalization of Markov Decision Processes [Puterman 1994] and as such are very different from the offline optimization setting for agent actions we consider here.

In a related tangent, the price of uncertainty [Balcan et al. 2009, 2011] studies the worst case implications to best response dynamics of small but adversarially selected fluctuations to the experienced cost of the agents. Results in this case tend to be negative. Carefully designed realizations of uncertainty combined with risk oblivious agents allow for small mistakes to pile up over time. In contrast, we show that even adversarial uncertainty of large magnitude can be effectively dealt away with by carefully chosen risk attitudes and approximately optimal global performance can be guaranteed.

Subsequently to our conference paper, Hota et al. [2013] analyze resource sharing games where a single resource comes with a probabilistic loss of investment for the agents due to resource failure. They model different risk attitudes (risk averse, risk neutral, and risk seeking) for the agents and show the existence of a pure strategy Nash equilibrium for certain instances of the game. They also give examples of inefficiency and convergence of best response dynamics. Angelidakis et al. [2013] also consider congestion games with risk sensitive agents. They focus on parallel-link networks and prove that such games 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. They also show some negative Price of Anarchy results for their model.

Lianeas [2014] studies when risk averseness of agents can reduce the price of anarchy in congestion game models with stochastic delays, and characterizes when a socially optimal traffic assignment can be implemented as a risk-averse equilibrium in series-parallel graphs. Nikolova and Stier-Moses [2014a] introduce the concept of the price of risk-aversion, defined as the ratio of the cost of risk-averse equilibrium to the cost of risk-neutral equilibrium. For homogeneous players with a mean-variance risk-averse objective, they prove a bound that is linear in the graph size and independent of the form of mean and variance functions of the edge delays. Meir and Parkes [2015b] study a congestion game, in which players have uncertainty over which routes other players use and take a non-probabilistic approach considering a specified range of players that use each edge. This is closer to the robust optimization framework in that agents minimize their worst-case cost. In [Meir and Parkes 2015a] smoothness type of arguments are developed for instances where agents are “playing the wrong game” which can capture several settings of games under nonstandard risk attitudes.

3. SINGLE AGENT DECISION MAKING UNDER UNCERTAINTY

To describe the issues involved in game theoretic settings under uncertainty, we start with the simpler single agent optimization setting that corresponds to decision making under uncertainty. Our setting will be built upon a static optimization problem under uncertainty.

A problem of optimization under uncertainty, which we will denote by $(\Gamma; \Omega; c; p)$, is specified by a feasible set Γ , a (possibly unknown) distribution p over the state of

nature, ω , chosen from set Ω , and an objective cost function $c : \Gamma \times \Omega \rightarrow \mathbb{R}$. We will assume this set Γ to be finite.

Over the last 60 years, work in decision theory has identified and formalized several different behavioral approaches to such uncertain environments. They are typically partitioned into the following two classes: a) *stochastic optimization*, that utilizes the distribution p over the uncertainty set and in some sense provides an “average-case” analysis, and b) *robust optimization*, whose guarantees hold even under worst case realizations of state uncertainty. Next, we will review some of these approaches in more detail.

3.1. Stochastic optimization

Robust optimization offers deterministic guarantees about the quality of the chosen decision. This power, however, comes at a cost since its proposed solutions may be overly conservative. Stochastic optimization utilizes extra information about the state uncertainty (e.g., a known probability distribution over the system states) to identify decisions of high *average-case* performance. Within this framework, several alternative decision-making principles have been identified.

Minimizing expected cost (risk-neutral agents)

In the case where the probability distribution p over the system states Ω is known, or at least can be estimated (e.g., by using historical data and/or a priori information), the optimal action for a risk neutral optimizer is the one that minimizes her expected cost. Formally, this corresponds to solving the following stochastic optimization problem:

$$\min_{\gamma \in \Gamma} \mathbb{E}_{\omega \sim p} [c(\gamma, \omega)]$$

Average case analysis

A less taxing computational heuristic for identifying solutions with good-on-average performance guarantees returns the optimal solution for the optimization problem in the expected state of the system. Namely, the random variable ω is replaced by its mean $\mathbb{E}_{\omega \sim p}[\omega]$, leading to the following deterministic optimization problem:

$$\min_{\gamma \in \Gamma} c(\gamma, \mathbb{E}_{\omega \sim p}[\omega])$$

The resulting optimal solution x is sometimes called an average case solution or an expected value solution. This technique is not always applicable⁵ and, even if it is, it does not necessarily translate to meaningful performance guarantees. Nevertheless, computationally it is rather inexpensive and hence it allows for a “quick fix” approach against which the performance of more sophisticated methods can be judged.

Win-or-Go-Home approach

This approach also optimizes the expected value of a specific random variable. Here, the random variable is an indicator variable that our chosen action outperforms another benchmark action (e.g., an action chosen by an opponent). Specifically, each agent (e.g., the i -th) is trying to maximize the probability of winning:

$$\Pr_{\omega \sim p} [\forall j \neq i, c_i(\gamma_i, \omega) \leq c_j(\gamma_j, \omega)]$$

⁵It is not always the case that the notion of expected system state is meaningful or that the cost function actually extends to these states.

Table II. Decision Making under Uncertainty

Principle	Implied Cost Function
Minimizing Expected cost	$\mathbb{E}_{\omega \sim p}[c(\gamma, \omega)]$
Average case analysis	$c(\gamma, \mathbb{E}_{\omega \sim p}[\omega])$
Win-or-Go-Home	$1 - \Pr_{\omega \sim p}[\forall o \in \text{Opponents}, c(\gamma, \omega) \leq c_o(\gamma_o, \omega)]$
Second moment method	$\mathbb{E}_{\omega \sim p}[c(\gamma, \omega)] + \alpha \cdot \text{Var}_{\omega \sim p}[c(\gamma, \omega)]$
Wald's minimax cost	$\max_{\omega \in \Omega} c(\gamma, \omega)$
Savage's minimax regret	$\max_{\omega \in \Omega} (c(\gamma, \omega) - \min_{\gamma' \in \Gamma} c(\gamma', \omega))$

In the case of ties above, we break them in favor of our chosen action. A less biased rule would choose a winner uniformly at random. Generally, such rules can be extended using numerous tie-breaking rules (*e.g.*, lexicographically). This line of thinking can also be trivially extended to include more than one opponent in which case we are trying to maximize the probability of outperforming all. This approach is especially well suited for competitive optimization settings under uncertainty [Immorlica et al. 2011] where the optimizers do not care about their actual performance but about outperforming their opponents (*e.g.*, heads-up contest, coming first in a race, etc.).

Incorporating second moments

So far, the decision-maker tries to optimize some cost function on average. However, the realized cost $c(\gamma, \omega)$ could be significantly different from the corresponding expected value $\mathbb{E}_{\omega \sim p}[c(\gamma, \omega)]$, depending on the particular realization of the system state ω . Considering $c(\gamma, \omega)$ as a random variable, we can hedge against such events by incorporating in our calculation the variance $\text{Var}[(\gamma, \omega)]$ of the action's cost.

$$\min_{\gamma \in \Gamma} \left(\mathbb{E}_{\omega \sim p}[c(\gamma, \omega)] + \alpha \cdot \text{Var}_{\omega \sim p}[c(\gamma, \omega)] \right)$$

The coefficient $\alpha \geq 0$ represents the weight given to the conservative part of the decision. As α grows to infinity, the above optimization problem tries to find a solution with minimal variance, while if $\alpha = 0$, this problem reduces to the standard risk neutral cost optimization. Note that since the variance is itself an expected value, from a mathematical point of view these problems still fall under the expectation minimization perspective. Other approaches could incorporate the standard deviation in place of the variance.

3.2. Robust optimization

Robust optimization techniques can be applied to any uncertain setting, regardless of having any distributional knowledge about the nature of the uncertainty. The recurring idea here is that we will always assume that the sources of uncertainty (*e.g.*, nature) act adversarially and once we commit to an action, the worst possible scenario (for that action) will be realized.

Wald's minimax (cost) principle

Wald's minimax principle [Wald 1939] is a deterministic decision-making model according to which decisions are ranked on the basis of their worst-case outcomes. Namely, the chosen decision is the one whose worst outcome is at least as good as the worst outcome of any other decision. Formally, Wald's minimax approach corresponds to solving the following deterministic problem:

$$\min_{\gamma \in \Gamma} \max_{\omega \in \Omega} c(\gamma, \omega).$$

This model represents a 2-person game in which the min optimization player plays first. In response, the second player (nature) selects the worst state in Ω , namely a state that maximizes the cost $c(\gamma, \omega)$. The above model is the “classic” format of Wald’s maximin model. There is an equivalent optimization or mathematical programming format:

$$\min_{\gamma \in \Gamma, z \in \mathbb{R}} \{z : z \geq c(\gamma, \omega), \forall \omega \in \Omega\}.$$

Essentially, any decision γ is assigned an (implied) cost $\max_{\omega \in \Omega} c(\gamma, \omega)$. This is usually referred to as “the security level” of decision γ .

Savage’s minimax regret principle

Savage’s minimax regret principle [Savage 1951] is an application of Wald’s minimax model to the ‘regret’ associated with the agent decision. It can be formulated as follows:

$$\min_{\gamma \in \Gamma} \max_{\omega \in \Omega} \left(c(\gamma, \omega) - \min_{\gamma' \in S} c(\gamma', \omega) \right)$$

where the term in the parenthesis captures the regret $r(\gamma, \omega)$ of the cost function c associated with the (decision, state) pair (γ, ω) .

4. CONGESTION GAMES WITH SCHEDULING POLICIES

4.1. Congestion games

Congestion games [Rosenthal 1973] are non-cooperative games in which the utility of each agent depends only on the agent’s strategy and the number of other agents that either choose the same strategy, or some strategy that “overlaps” with it. Formally, a *congestion game* is defined by the tuple $(N; E; (S_i)_{i \in N}; (c_e)_{e \in E})$ where N is the set of *agents*, E is a set of *resources* (also known as *edges*, *bins*, or *facilities*), and each agent i has a set S_i of subsets of E ($S_i \subseteq 2^E$). Each strategy $s_i \in S_i$ is a set of edges (*a path*), and c_e is a cost (negative utility) function associated with edge e . In symmetric congestion games, all agents share the same strategy set, *i.e.*, $S_i = S_j$ for all i, j . For a strategy profile $s = (s_1, s_2, \dots, s_N)$, the cost of agent i is given by $c_i(s) = \sum_{e \in s_i} c_e(\ell_e(s))$, where $\ell_e(s)$ is the number of agents using e in s (the load of edge e). Congestion games admit a *potential function* $\Phi(s) = \sum_{e \in E} \sum_{j=1}^{\ell_e(s)} c_e(j)$, which captures each agent’s incentive to change her strategy [Rosenthal 1973]. Specifically, given a strategy profile $s = (s_1, s_2, \dots, s_N)$, and strategy s'_i of agent i , we have $c_i(s'_i, s_{-i}) - c_i(s) = \Phi(s'_i, s_{-i}) - \Phi(s)$. As a result, starting from any strategy profile, any sequence of improving moves is bound to terminate. Such stable states $s \in S$ where for each agent i and $s'_i \in S_i$, $c_i(s'_i, s_{-i}) \geq c_i(s)$ are called Nash equilibria. The set of Nash equilibria can be found by simply locating the local optima of $\Phi(S)$. In this work, we will focus primarily on linear congestion games where the latency functions are of the form $c_e(x) = a_e x + b_e$ where $a_e, b_e \geq 0$ and x is the number of users on edge e . Similarly, in polynomial congestion games, the latency functions are polynomial functions with nonnegative coefficients. Furthermore, the social cost will correspond to the sum of the costs of all the agents $\sum_i c_i(s)$. The price of anarchy is defined as:

$$\text{PoA}(G) = \frac{\max_{s \in \text{NE}} \text{Social-Cost}(s)}{\min_{s^* \in \times_i S_i} \text{Social-Cost}(s^*)}.$$

4.2. Congestion games with randomizing schedulers

We will begin the discussion with an informal, high level description of our model. In a classic congestion game all agents using a specific resource experience the same latency as is defined by the total number of the agents using a resource. Here, we will assume that the agents using any machine are assigned an ordering (*e.g.*, they enter a queue) and the latency of an agent only depends on the number of other agents that precede her in the queue. The ordering decisions on each individual machine are made according to some probability distribution. As a result, the rank of each agent on each machine (and therefore her respective experienced latency) is a random variable. This defines a strategic game with uncertain cost functions. Our goal will be to analyze the nature and efficiency of its stable states under different decision-theoretic approaches to environmental uncertainty.

More formally, we extend the standard congestion game setting by introducing a randomizing scheduler on each edge e . The scheduler chooses probabilistically a ranking (permutation) over all agents. We denote the rank of agent i using edge e by $r_e(i)$. This is a random variable taking values $\{1, 2, \dots, N\}$. Finally, the cost of agent i using a set of resources s can be expressed as $c_i(s) = \sum_{e \in s_i} c_e(\sum_{j: e \in s_j \wedge r_e(j) \leq r_e(i)} 1)$, since it is only affected by the number of agents that appear earlier than she on each edge $e \in S$.

4.3. Models of Randomizing Schedulers

For our core model the random rank of all agents is going to be uniformly chosen (between 1 and N) and identical across all congestion elements/edges. Namely, we choose a single uniformly random permutation of all agents and their respective order on each edge that they use respects this ordering. For example if agents i, j follow the same route then if agent i is routed before agent j along their first edge then they will keep that respective ordering throughout their path. Such phenomena arise in practice in real-life routing situations. In this sense, this model can be thought of as a first order approximation of capturing traffic flow issues within the stylized setting of congestion games. Of course, this model by no means perfectly captures all flow constraints but it is still an improvement over the standard congestion game/model where the ordering of the agents is not even accounted for.

Another application that such models can capture is the possibility of tiered systems where all users of type α (*e.g.*, gold members) get preferential treatment over all users of type β (*e.g.*, general population). Our model captures an extremal such case where each individual corresponds to a type and where the user is unaware of where she stands in the global pecking order. As we show, the price of anarchy is improved in comparison to the standard linear congestion games and this may indicate that such techniques may actually improve the overall efficiency of routing systems. Validating this claim, however, would require a careful analysis of the effects of less fine-grained type models which we leave as an open challenge for future work.

Variants: All of our results carry over even if the uniformly random orderings are independent across different edges. Such a model might be a better approximation of scheduling in load balancing settings. For example, if there is a series of printers and two (or more) jobs are allocated to the same machine at the same time then they will be executed consecutively in effectively random order. In this case it makes more sense to think of the random scheduling of each machine to be independent of each other. Moreover, the results about risk neutral optimization as well as the average case analysis extend (due to linearity of expectation) to the case where the random orderings of different schedulers are correlated, as long as the marginal distribution of each scheduler assigns each agent to each slot with equal probability ($1/N$). A case of particular interest is the case of Wald's minimax cost. In this case, the price of

anarchy analysis holds for an even more expansive class of randomizing schedulers. Specifically, we allow any set of (possibly adversarially) correlated random orderings as long as for any agent i there is a positive probability that she will be ordered last in all resources.

4.4. Social Cost and Efficiency

We apply each of the different decision-making principles in our setting of congestion games with randomizing schedulers. Each principle, when combined with its respective implied cost functions (Table II), induces novel deterministic classes of games. We will examine the nature of the Nash equilibria of each resulting class and provide tight characterizations about their price of anarchy (Table I). In order to do so, we need to give a definition of social cost for each outcome. The social cost of any strategy outcome will be estimated according to the experienced (real-world) costs of the outcomes and not the implied costs of Table II. After all, the implied costs are merely a mathematical convenience for encoding the decision making process of risk sensitive users in uncertain environments. These elements are out of the control and possibly the knowledge of the system designer. The system designer instead cares about the realized system performance and about achieving near optimality despite lack of control over both agents' actions and risk attitudes.

Formally, we define the social cost in its standard form as the sum of the (real world) costs of the agents. These individual costs of any user on a resource are random variables that depend on her ordering on that resource. On the other hand, the sum of all users' costs (*i.e.*, the social cost of any outcome) is deterministic. The social cost is in effect equal to $\sum_{e \in E} \sum_{j=1}^{\ell_e(s)} c_e(j)$, which coincides with the definition of potential function for the standard formulation of congestion games. This definition holds regardless of the risk attitudes of the agents. This exactly captures the sum of the actual cost of all agents since on each element that experiences load x there exists exactly one agent placed first, another placed second and so on. Naturally, from the perspective of the system designer a small price of anarchy is highly desirable since it implies approximately optimal system performance at equilibrium.⁶

Although there are several ways of introducing uncertainty in a congestion game, our approach is especially well suited to study the effects of different decision-theoretic principles on the efficiency of these games. Our setting allows one to study the effects of different optimization approaches in isolation of other factors influencing the price of anarchy. Indeed, since the social cost for any outcome is deterministic regardless of the chosen decision-making process, all of these models agree in terms of which are the socially optimal states. Furthermore, any two of them assign the same social cost to any state, not just the optimal ones. As a result, changes in the value of the price of anarchy can only be introduced by having different sets of equilibrium states corresponding to different optimization approaches. For example, an outcome that may be stable for risk neutral agents might not be stable for agents that apply robust optimization techniques, or the other way around. Naturally, the question that arises is how robust is the price of anarchy analysis to changes in the risk attitude of the agents? Table I shows our results in each case.

⁶The system designer, as an optimizer, may also have its own attitude towards risk, which could be different from the ones of the agents participating in the system. Since in our model of uncertainty the social cost of each state is deterministic, our price of anarchy bounds hold independently of the risk attitude of the system designer.

5. PRICE OF ANARCHY UNDER RISK NEUTRAL STOCHASTIC OPTIMIZATION

We will commence the analysis with the, thus far, standard assumption in algorithmic game theory of risk neutral agents. The agents here merely wish to minimize their expected costs. Hence, and in accordance with Table II, we have that the implied cost \mathcal{C}_i of agent i at strategy profile s is $\mathbb{E}_{\omega \sim p}[c_i(s, \omega)]$.

We will assume our standard model of randomizing schedulers where all schedulers process the users according to a common and uniformly randomly chosen permutation. Moreover, these results can still be straightforwardly generalized to other models of randomizing schedulers as discussed in section 4.2 due to linearity of expectation as long as the marginal distribution of each scheduler assigns each agent to each slot with equal probability ($1/N$).

LEMMA 5.1. *The price of anarchy of a linear congestion game $(N; E; (S_i)_{i \in N}; (c_e)_{e \in E})$ with uniformly randomizing schedulers under risk neutral agents is equal to the price of anarchy of the linear congestion game $(N; E; (S_i)_{i \in N}; (c'_e)_{e \in E})$ where $c'_e(x) = c_e(\frac{x+1}{2})$.*

PROOF. Under this assumption, for linear cost functions the implied cost of an agent i at outcome s is:

$$\begin{aligned} \mathcal{C}_i(s) &= \mathbb{E}_{\omega \sim p}[c_i(s, \omega)] = \sum_{e \in s_i} \mathbb{E}_{\omega \sim p} \left[c_e \left(\sum_{j: e \in s_j \wedge r_e(j) \leq r_e(i)} 1 \right) \right] = \sum_{e \in s_i} \frac{\ell_e(s) \left(a_e \frac{\ell_e(s)+1}{2} + b_e \right)}{\ell_e(s)} \\ &= \sum_{e \in s_i} \left(a_e \frac{\ell_e(s)+1}{2} + b_e \right) = \sum_{e \in s_i} \left(\frac{a_e}{2} \ell_e(s) + \frac{a_e}{2} + b_e \right) \end{aligned}$$

It is straightforward to check that starting from any congestion game with affine cost functions (*i.e.*, $a_e, b_e \geq 0$ for all resources e), this transformation through the implied costs $\mathcal{C}_i(s)$ maps us back in the same class. Furthermore, at any strategy profile s (and for any possible realization of the uncertainty ω) the actual social cost agrees with the implied social cost, defined as the sum of implied user costs $\sum_i \mathcal{C}_i(s) = \sum_i c_i(s, \omega)$ for all s, ω . \square

COROLLARY 5.2. *The price of anarchy for risk neutral agents under affine cost functions is equal to the price of anarchy of the class of congestion games with cost functions of the form $a_e(x+1) + b_e$ where $a_e, b_e \geq 0$. Equivalently, it suffices to characterize those congestion games whose elements have cost functions $a_e x + b_e$ with $b_e \geq a_e \geq 0$.*

Since this class defines a subset of affine congestion games, we automatically derive an upper bound on the price of anarchy of $5/2$ [Christodoulou and Koutsoupias 2005]. A more thorough analysis will allow us to improve this estimate by 33% and prove a tight bound of $5/3$. It is helpful to keep in mind that $\mathcal{C}(s) = \sum_e (a \ell_e(s) + b) \ell_e(s)$ and $\Phi(s) = \sum_e \sum_{j=1}^{\ell_e(s)} c_e(j) = \sum_e a_e \frac{\ell_e(s)(\ell_e(s)+1)}{2} + b \ell_e(s)$. As we state in section 4.4, Φ captures the social cost in the case of our congestion games with schedulers. All the arguments in this section critically depend on the assumption of linear latency functions.

Next, we will prove a smoothness bound argument about our class of congestion games. A cost-minimization game is (λ, μ) -smooth [Roughgarden 2009] if for every two outcomes s and s^* , $\sum_i c_i(s_i^*, s_{-i}) \leq \lambda \cdot \text{Social-Cost}(s^*) + \mu \cdot \text{Social-Cost}(s)$, which if $\lambda > 0$ and $\mu < 1$ can be shown to imply a $\frac{\lambda}{1-\mu}$ upper bound on the price of anarchy. In the following, $\mathcal{C}(s)$ denotes the implied social cost at outcome s , that is $\sum_i \mathcal{C}_i(s)$.

THEOREM 5.3. *Given any two strategy profiles s, s^* of a linear congestion game with uniformly randomizing schedulers, under risk-neutral agents, we have that:*

$$\sum_i C_i(s_i^*, s_{-i}) \leq \frac{5}{4}C(s^*) + \frac{1}{4}C(s).$$

PROOF. Let $x_e = \ell_e(s), y_e = \ell_e(s^*)$. Then,

$$\begin{aligned} \sum_i C_i(s_i^*, s_{-i}) &\leq \sum_e (a_e(x_e + 2) + b_e)y_e \\ &\leq \frac{5}{4} \sum_e (a_e(y_e + 1) + b_e)y_e + \frac{1}{4} \sum_e (a_e(x_e + 1) + b_e)x_e \\ &\leq \frac{5}{4}C(s^*) + \frac{1}{4}C(s). \end{aligned}$$

The second inequality follows from lemma A.1 in appendix A. \square

Theorem 5.3 expresses a $(5/4, 1/4)$ -smoothness argument for our class of games, which by [Roughgarden 2009] implies that our class of games exhibits price of anarchy of $5/3$.

COROLLARY 5.4. *The price of anarchy of a linear congestion game with uniformly randomizing schedulers under risk-neutral agents is at most $5/3$.*

Next, we will show that our price of anarchy bound is tight for the class of linear congestion games. Furthermore, it is asymptotically tight even for symmetric congestion games under risk-neutral stochastic optimization. Below we provide explicit constructions of classes of games achieving these lower bounds. These constructions are inspired by the work of Christodoulou and Koutsoupias [2005].

THEOREM 5.5. *There exist linear congestion games whose price of anarchy under uniformly randomizing schedulers and risk-neutral agents is equal to $5/3$. Furthermore, for any $\epsilon > 0$, there exist symmetric linear congestion games whose price of anarchy under uniformly randomizing schedulers and risk-neutral agents is greater than $\frac{5}{3} - \epsilon$.*

PROOF. We will construct a congestion game for $N \geq 3$ players and $|E| = 2N$ resources with price of anarchy equal to $5/3$. We divide the set E into two subsets $E_1 = \{h_1, \dots, h_N\}$ and $E_2 = \{g_1, \dots, g_N\}$, each of N resources. Player i has two pure strategies: $\{h_i, g_i\}$ and $\{h_{i+1}, g_{i-1}, g_{i+1}\}$ ⁷. The latency functions of the elements in E_1 are $c(x) = 2x$, whereas the latency functions of the elements in E_2 are $c(x) = x$. The optimal outcome corresponds to the state where each player selects the first strategy while the worst-case Nash equilibrium is for each player to select the second strategy. The second outcome is a Nash, since under uniformly random scheduling each agent experiences an expected cost of $2 + 2\frac{1+2}{2} = 5$, whereas if she were to deviate she would experience an expected cost of $\frac{2+4}{2} + \frac{1+2+3}{3} = 5$. At the worst Nash equilibrium, the total expected cost is equal to $5N$. At the optimal state, the expected cost is equal to $3N$, which results in a ratio of $5/3$.

We will show how to construct symmetric linear congestion games with price of anarchy arbitrary close to $5/3$ under risk neutral agents. It suffices to show that we can construct a classic congestion game, in which each element has cost function $x + 1$, and

⁷The indices i should appear as $i \bmod N$, but we suppress it for simplicity of notation.

show that its price of anarchy can be arbitrarily close to $5/3^8$. We construct such games as follows: We partition the resources into sets $\{P_1, P_2, \dots, P_n\}$ of the same cardinality and make each P_i a pure strategy. Each P_i contains $N\alpha_1 + \binom{N}{2}\alpha_2$ resources where α_1, α_2 are appropriate constants to be determined later. We furthermore partition each of the $N\alpha_1$ resources of set P_i into N groups $\{R_1^{P_i}, \dots, R_N^{P_i}\}$, each of size α_1 . Similarly, we partition the $\binom{N}{2}\alpha_2$ resources of set P_i into $\binom{N}{2}$ groups $Q \{Q_{12}^{P_i}, \dots, Q_{ij}^{P_i}, \dots, Q_{N-1N}^{P_i}\}$ each of size α_2 .

The strategies available to each player i will include, besides all the P_i strategies, strategies (A_1, \dots, A_n) of the following form: Each strategy A_i contains out of each P_j , all the resources in $R_i^{P_j}$ as well as all resources in each of the $Q_{kl}^{P_j}$ where either $k = i$ or $l = i$. As a result, each strategy A_i includes $N\alpha_1 + N(N-1)\alpha_2$ elements.

The strategy profile (P_1, P_2, \dots, P_n) is trivially a Nash equilibrium of the game, since each P_i is an action with a minimal number of resources (amongst all available actions) and the load of each facility is optimal (*i.e.*, equal to 1). It is also the social optimal outcome for the same reason, with a cost for each agent of $2N\alpha_1 + N(N-1)\alpha_2$.

We choose α_1, α_2 , so that the strategy profile (A_1, A_2, \dots, A_n) is also a Nash equilibrium. Now, the cost for each player i is $c_i(A) = N(2\alpha_1 + 3(N-1)\alpha_2)$. It suffices to select α_1, α_2 so that player i will not switch to P_j . (It is straightforward that player i will not switch to the Nash strategy of some other player k .) The cost after switching is:

$$c_i(A_{-i}, P_j) = 2\alpha_1 + 3(N-1)\alpha_2 + 3(N-1)\alpha_1 + 4\binom{N-1}{2}\alpha_2 = (3N-1)\alpha_1 + (N-1)(2N-1)\alpha_2$$

To establish Nash equilibrium it suffices to have $c_i(A) = c_i(A_{-i}, P_j)$, or equivalently $\alpha_1 = (N+1)\alpha_2$, which is satisfied when we select $\alpha_1 = N+1$ and $\alpha_2 = 1$. With this, the social cost at the (A_1, A_2, \dots, A_n) Nash equilibrium is $N^2(2\alpha_1 + 3(N-1)\alpha_2) = N^2(5N-1)$. The social welfare at the (P_1, P_2, \dots, P_n) Nash equilibrium is equal to $N^2(3N+1)$. Hence, as N goes to infinity, their ratio converges to $5/3$. \square

Combining all results in this section, we derive a complete understanding of the effects of risk neutral optimization under uniformly randomizing agents in linear congestion games:

THEOREM 5.6. *Given any linear congestion game with arbitrary topology and uniformly randomizing schedulers, the price of anarchy under risk neutral stochastic optimization is at most $5/3$ and this bound is tight.*

6. PRICE OF ANARCHY UNDER AVERAGE CASE ANALYSIS

Given any outcome s of a congestion game with randomizing schedulers, we can encode the randomness relevant to agent i via a sequence of variables $r_e^i(s) \in \{1, \dots, \ell_e(s)\}$, which encodes the effective order with which agent i 's job is processed by machine e . On average, the order of agent i on machine e is merely the expected rank $\hat{r}_e^i(s) = \mathbb{E}_{\omega \sim p}[r_e^i(s)]$. Naturally, this expectation is not always an integer, so the average case analysis approach is not immediately applicable in our setting. Namely, if we consider two possible outcomes s, s^* , the superposition of states that corresponds to s (with probability p) and s^* (with probability $1-p$) is not a meaningful state. However, one could extend the setting in a straightforward way so as to include such ‘‘mixed’’ states. In such a state, the load of each machine is by definition equal to expected load of the machine with the expectation taken over the pure states of the system. Finally, the

⁸Indeed, this game would have the same price of anarchy as the corresponding congestion game with edge cost $2x$ under uniform schedulers and risk neutral agents.

(linear) cost functions can be augmented so that they are defined over all nonnegative reals (instead of merely nonnegative integers) while preserving the linearity.

LEMMA 6.1. *The price of anarchy of a linear congestion game $(N; E; (\mathcal{S}_i)_{i \in N}; (c_e)_{e \in E})$ with uniformly randomizing schedulers under average case analysis is equal to the price of anarchy of the linear congestion game $(N; E; (\mathcal{S}_i)_{i \in N}; (c'_e)_{e \in E})$ where $c'_e(x) = c_e(\frac{x+1}{2})$.*

PROOF. The implied cost functions under average case analysis are as follows:

$$C_i(s) = c_i\left(s, \mathbb{E}_{\omega \sim p}[\omega]\right) = \sum_{e \in s_i} c_e\left(\mathbb{E}_{\omega \sim p}[\ell_e(s)]\right) = \sum_{e \in s_i} c_e\left(\frac{\ell_e(s) + 1}{2}\right).$$

□

However, this game is completely isomorphic to the one arising from the case of risk-neutral agents, so our tight results in that section get immediately translated here:

PROPOSITION 6.2. *Given any linear congestion game with arbitrary topology and uniformly randomizing schedulers, the price of anarchy under the average case analysis model is at most $5/3$ and this bound is tight.*

7. PRICE OF ANARCHY UNDER WIN-OR-GO-HOME PRINCIPLE

Our results under different risk attitudes have been universally positive so far and actually improving upon the classic $5/2$ price of anarchy bound of linear congestion games. Next, we will show that different approaches to uncertainty can actually have a detrimental effect on system performance.

Recall that in this model, we are trying to maximize the probability of outperforming a benchmark action (e.g., an opponent). Specifically, each agent (e.g., the i -th) is trying to maximize the probability of winning:

$$\Pr_{\omega \sim p}[\forall j \neq i, c_i(\gamma_i, \omega) \leq c_j(\gamma_j, \omega)]$$

In our setting, we can consider all other agents as a single opponent, in which case we wish to maximize the probability of experiencing the lowest cost amongst all agents (i.e., arriving first to our destination).

In the case of ties, we break ties in favor of our chosen action. A less biased rule could choose a winner uniformly at random. Generally, such rules can be extended using numerous tie-breaking rules (e.g., lexicographically). Our results trivially carry over to these models as well.

PROPOSITION 7.1. *Under uniformly randomizing schedulers, for any $M \geq 0$, there exist linear congestion games such that under the Win-or-Go-Home principle the price of anarchy is greater than M .*

PROOF. Let's assume a load balancing game with N agents and N resources $\{1, 2, \dots, N\}$ with cost functions $c_1(x) = x$ and $c_i(x) = x + 1$. The strategy outcome where all agents choose the first resource is a Nash equilibrium. Indeed, the probability of any agent of having the minimal cost amongst all agents is $1/N$. If an agent deviates then his cost would be equal to 2 with probability 1, however there will always be someone outperforming him for any possible ordering of the rest of the agents on the first resource. Hence, the probability of "winning" drops to zero when deviating to another strategy. Finally, the social cost of this strategy is $N(N + 1)/2$, whereas the optimal is $2N - 1$. The price of anarchy for the resulting class of games in unbounded and the theorem follows. □

8. PRICE OF ANARCHY UNDER SECOND MOMENT STOCHASTIC OPTIMIZATION

We will show that under this risk attitude the price of anarchy explodes once again.

PROPOSITION 8.1. *Under uniformly randomizing schedulers, for any $M \geq 0$, there exist linear congestion games with only two paths, and $\alpha \geq 0$, such that under α -stochastic optimization, their price of anarchy is greater than M .*

PROOF. Consider a congestion game of N agents consisting of a network of two parallel edges, one with a constant latency of L and another with a linear latency function $c_e(x) = x$. If the load of the second edge is greater than two, then the variance of the cost of the edge for each agent would be strictly positive as a result of the random ordering of the agents. On the other hand, the variance of the cost of each agent on the first machine is always zero, since the latency is load (and hence ordering) independent. Therefore, for any N, L , we can choose a α large enough so that any α -stochastic optimization agent (whose best response is $\min_{s \in S} (\mathbb{E}_{\omega \sim p} [c_i(s, \omega)] + \alpha \cdot \text{Var}_{\omega \sim p} [c_i(s, \omega)])$) prefers using the first edge as long as the second edge is occupied by at least one agent. We will argue that for large enough α the strategy outcome with exactly one agent using the linear latency resource is a Nash equilibrium. Indeed, under a uniformly random scheduling policy and a load of x , the variance of the linear resource will be equal to $\frac{x^2-1}{12}$. So, for $x \geq 2$ the variance is at least $1/4$. Hence, as long as $\alpha > 4L$, the agents using the first resource (experiencing a cost of L) do not benefit by migrating to the second resource. The single agent on the second resource (experiencing a cost of 1) similarly does not wish to deviate to the first resource. By picking $L = \Omega(N^2)$, we have that the price of anarchy of the resulting class of games is unbounded and the theorem follows. \square

9. PRICE OF ANARCHY UNDER WALD'S MINIMAX PRINCIPLE

In this section, we will provide a price of anarchy analysis in the case where agents apply robust optimization techniques and specifically Wald's minimax principle. Following this principle, routes are ranked on the basis of their worst-case (scheduling) outcomes. In other words, each agent chooses the route with the minimal worst case latency. As long as each scheduler assigns to each agent i each slot with positive probability (we will refer to such schedulers as fully randomizing schedulers), the implied cost of agent i is equal to:

$$C_i(s) = \max_{\omega \in \Omega} c_i(s, \omega) = \sum_{e \in s_i} c_e(\ell_e(s)) = \sum_{e \in s_i} (a_e \ell_e(s) + b_e)$$

Namely, we have that this implied-cost game is payoff equivalent to the classic formulation of (linear) congestion games. However, regardless of the decision making process of the agents and its depiction through the implied costs $C_i(s)$, the actual social cost at any outcome s is equal to:

$$\text{Social-Cost}(s) = \sum_e \sum_{k=1}^{\ell_e(s)} c_e(k) = \sum_e \sum_{k=1}^{\ell_e(s)} (a_e k + b_e)$$

This is true since on each resource e agents are only burdened by the load corresponding to other agents who are ahead of them in the queue. Furthermore, although the exact ordering and costs of each agent are unclear, there will always be someone occupying the first, second, . . . , k -th spot. So, the summation of the costs is indeed captured by the summation above. This summation encodes the standard potential function for (linear) congestion games. In order to characterize the price of anarchy under

Wald's minimax principle, it suffices to analyze the classic setting when we define the social cost to be equal to the potential function $\sum_e \sum_{k=1}^{\ell_e(s)} c_e(k)$. Independently from us, Caragiannis et al. [2011] provide an upper bound of 2 for this term in their approach to efficiently compute approximate Nash equilibria in linear congestion games. We provide an alternative proof of this statement, which tracks the involved summands more closely. Our technique, which is based on technical lemmas B.1–B.2 in the appendix, can also be applied to the case of Savage's minimax principle to provide non-tight upper bounds on the price of anarchy. In contrast, in the case of Wald's minimax principle, our price of anarchy bounds are tight.

We prove an upper bound of 2 for the ratio between the value of the potential function $\Phi(s) = \sum_e \sum_{k=1}^{\ell_e(s)} c_e(k) = \sum_e \sum_{k=1}^{\ell_e(s)} (a_e k + b_e)$ at the worst Nash equilibrium and its optimum value.⁹

THEOREM 9.1. *The value of the potential function $\Phi(s) = \sum_e \sum_{k=1}^{\ell_e(s)} c_e(k) = \sum_e \sum_{k=1}^{\ell_e(s)} (a_e k + b_e)$ of any linear congestion game at any Nash equilibrium is within a factor of two of its optimal value.*

PROOF. Suppose that (s_1, s_2, \dots, s_n) is a Nash equilibrium of a linear congestion game and $(s_1^*, s_2^*, \dots, s_n^*)$ corresponds to the global optimum of the potential function. Then:

$$\begin{aligned}
\Phi(s_1, s_2, \dots, s_n) &= \sum_e \sum_{i=1}^{x_e} (a_e i + b_e) = \sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e b_e x_e \\
&= \frac{1}{2} \sum_e (a_e x_e + b_e) x_e + \frac{1}{2} \sum_e (a_e + b_e) x_e = \frac{1}{2} \sum_i \mathcal{C}_i(s) + \frac{1}{2} \sum_e (a_e + b_e) x_e \\
&\leq \frac{1}{2} \sum_i \mathcal{C}_i(s_i^*, s_{-i}) + \frac{1}{2} \sum_e (a_e + b_e) x_e \leq \frac{1}{2} \sum_e (a_e(x_e + 1) + b_e) x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e \\
&= \frac{1}{2} \sum_e a_e x_e^* (x_e + 1) + \frac{1}{2} \sum_e b_e x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e \\
&= \frac{1}{2} \left(\frac{1}{2} \sum_e a_e \frac{x_e(x_e + 1)}{2} + 2 \sum_e a_e \frac{x_e^*(x_e^* + 1)}{2} - \frac{1}{4} \sum_e a_e x_e - \frac{1}{4} \sum_e a_e (x_e - 2x_e^*)^2 \right) \\
&+ \frac{1}{2} \sum_e b_e x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e \text{ from Lemma B.1} \\
&= \frac{1}{4} \sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e a_e \frac{x_e^*(x_e^* + 1)}{2} + \frac{1}{2} \sum_e b_e x_e^* - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 + \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + b_e \right) x_e \\
&\leq \frac{1}{4} \Phi(s_1, s_2, \dots, s_n) + \Phi(s_1^*, s_2^*, \dots, s_n^*) - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 + \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + \frac{b_e}{2} \right) x_e.
\end{aligned}$$

By applying Lemma B.2, we have that:

$$\Phi(s_1, s_2, \dots, s_n) \leq \frac{1}{2} \Phi(s_1, s_2, \dots, s_n) + \Phi(s_1^*, s_2^*, \dots, s_n^*)$$

The statement of the theorem follows immediately. \square

⁹The optimal value is also achieved at a Nash equilibrium, since it is the global minimizer of the potential.

We also establish that this bound is indeed the best possible. The examples of games that provide the matching lower bounds are similar to the ones providing the tight bounds for risk neutral agents, however in this case we are focusing on the behavior of the potential function Φ .

THEOREM 9.2. *There exist linear congestion games such that the maximum ratio of the value of the potential function $\Phi(s) = \sum_e \sum_{k=1}^{\ell_e(s)} c_e(k) = \sum_e \sum_{k=1}^{\ell_e(s)} (a_e \ell_e(s) + b_e)$ at any Nash equilibrium and its optimum value is equal to 2. Furthermore, for any $\epsilon > 0$, there exist linear symmetric congestion games such that the corresponding ratio is greater than $2 - \epsilon$.*

Combining all the theorems of this section together, we finally derive a tight price of anarchy analysis in the case of agents applying Wald's minimax principle:

THEOREM 9.3. *Given any linear congestion game with arbitrary topology, and any set of fully randomizing schedulers, (even adversarial, correlated), the price of anarchy under robust optimization is at most 2 and this bound is tight.*

10. PRICE OF ANARCHY UNDER SAVAGE'S MINIMAX REGRET PRINCIPLE

In this section we will present some preliminary results about price of anarchy in linear symmetric load balancing games under Savage's minimax regret principle. Load balancing games define a special class of congestion games where all strategies are singletons (i.e. paths of length 1). Finally, symmetric congestion games have the property that all agents share the same strategy set. Savage's minimax regret principle essentially applies Wald's minimax model to the regrets associated with each agent's decision. The special structure of symmetric load balancing games allows for a more concrete representation of each agent's estimated cost under Savage's principle. Specifically, the term $\max_{\omega \in \Omega} (c(s, \omega) - \min_{s' \in S} c(s', \omega))$ reduces to $\max\{c_e(\ell_e(s)) - \min_{e' \neq e} c_{e'}(1), 0\}$. The reason why we have to consider taking max over zero is that if the current latency of the agent for being assigned last to his current machine is less or equal to his latency for being assigned first to second faster machine, then the optimal action for the agent even under worst case realization of uncertainty is to choose his current machine. In terms of recognizing Nash equilibria it suffices to use the cost function $c_e(\ell_e(s)) - \min_{e' \neq e} c_{e'}(1)$.

LEMMA 10.1. *A state s in a load balancing game under independent, uniformly randomizing schedulers, and Savage's minimax principle is a Nash equilibrium if and only if it is a Nash equilibrium in a load balancing game with the same set of agents and strategies/edges where the latency function of each edge is equal to $C_e(s) = c_e(\ell_e(s)) - \min_{e' \neq e} c_{e'}(1)$.*

PROOF. Given any outcome $s = (e, s_{-i})$, since the schedulers are randomizing independently the worst case experienced cost for agent i on resource e is when he is assigned last on that element. Naturally, his latency in that case would be equal to $c_e(\ell_e(s))$. In terms of regret, since the schedulers are independent his regret is maximized when he considers that he is placed first in the random ordering of the machine that he chooses to deviate to. Naturally, the best such machine is the one with minimal $c_{e'}(1)$ amongst the remaining elements/machines. In the special case where $c_e(\ell_e(s)) - \min_{e' \neq e} c_{e'}(1) < 0$, the best action for the agent is to stay at his current strategy where he enjoys zero regret/cost. However, since the regrets of all other actions is nonnegative even if we assign $c_e(\ell_e(s)) - \min_{e' \neq e} c_{e'}(1)$ cost to this action we do not alter the set of equilibria of the game. \square

The above lemma immediately establishes that all such resulting games are still congestion games and therefore pure Nash equilibria always exist. It is easy to check that if there exist at least two machines with the same minimal latency under load of one job, then, for all agents i and outcomes s , we have $C_i(e, s_{-i}) = c_e(l_e(s)) - \min_e c_e(1)$. However, this game is payoff equivalent to the game arising under Wald's minimax principle resulting in a price of anarchy of 2. We will show that even when this hypothesis does not hold the price of anarchy is still bounded (at most equal to 3).

In the general case, let e_0 be the machine with the minimal latency under load of 1 and let e_1 be the machine that minimizes this latency amongst the rest. Let $\Delta = c_{e_1}(1) - c_{e_0}(1) \geq 0$ be the difference between the latencies of these two machines under load of one. The resulting game as described in Lemma 10.1 is equivalent to one with latency functions $c_{e_0}(x) = c_{e_0}(x)$ and $c_{e'}(x) = c_{e'}(x) + \Delta$ for all $e' \neq e_0$.

THEOREM 10.2. *The price of anarchy for any symmetric linear load balancing game under independent uniformly randomizing schedulers and agents applying Savage's minimax principle is at most 3.*

Next, we prove some lower bounds on the price of anarchy. We will show that even in this restricted case of symmetric linear load balancing game under uniform independently randomizing schedulers the price of anarchy is at least $4/3$ whereas if we extend our inquiries into linear congestion games then the price of anarchy can be greater or equal to $5/3$.

THEOREM 10.3. *The price of anarchy of symmetric linear load balancing games under independent uniformly randomizing schedulers and agents applying Savage's minimax principle is at least $4/3$. Furthermore, there exist linear (non load-balancing) congestion games with a price of anarchy of $5/3$.*

Putting everything together we derive the following statement.

COROLLARY 10.4. *The price of anarchy for any symmetric linear load balancing game under independent uniformly randomizing schedulers and agents applying Savage's minimax principle is at most 3. There exist symmetric linear load balancing games with a price of anarchy of $4/3$. Furthermore, there exist linear (non load-balancing) congestion games with a price of anarchy of $5/3$.*

11. CONCLUSIONS

In this paper we have introduced risk-sensitive agents in competitive environments under uncertainty and we have shown that the choice of risk attitudes can be critical to system performance. Specifically, our treatment allows us to isolate and examine the effects of a single parameter (risk attitude towards uncertain cost functions) on system performance. Since uncertainty is an inherent characteristic of decentralized systems and risk attitudes are out of the control of any system designer (and sometimes out of the conscious control of the agents themselves), our results call for a closer reexamination of competitive settings that were previously thought of as completely understood. More generally, a principled study of different sources of uncertainty and their implications to the performance and computational tractability of classic competitive settings would shed much needed light towards a deeper quantitative understanding of the implications of risk attitudes.

We choose here to focus on this single issue and show that this parameter single-handedly can have a critical effect on system performance. Nevertheless, we recognize that there is a whole host of related issues, which constitute interesting directions for future research. At the intellectual core of these issues lies the implicit assumption of a permanent, deterministic and context-free utility function that trivializes the task of

choosing between two (or more) alternatives for each agent. The problem of choice is inherently an algorithmic one and we believe that sidestepping the rigid assumptions of utility theory and focusing on it directly holds the promise of exciting new insights.

Our approach currently focuses mostly on the well studied, but arguably narrow, class of linear congestion games, however, we contrast this by examining a rather wide range of risk attitudes. Given the fact that linear congestion games have been a benchmark in the study of price of anarchy, they constitute a natural starting choice when trying to extend such results to more complicated, realistic risk attitudes. Furthermore, in appendix D we present (non-tight) upper bounds on the price of anarchy for polynomial cost functions. Tightening these upper bounds would be an interesting, nontrivial extension of our current work. It would be interesting to focus on other classes of games (e.g. non-potential games) as well as other sources of payoff uncertainty. Finally, it would be exciting to understand how uncertainty can affect the performance and topology of learning dynamics, e.g., [Papadimitriou and Piliouras 2016; Panageas and Piliouras 2016; Ligett and Piliouras 2011].

We believe that this intersection between algorithmic game theory and algorithmic decision theory including robust/stochastic optimization provides a rather promising field for future work. Here, algorithmic game theory can move beyond merely addressing the computational tractability of old game theoretic concepts but also help lay the foundational work for a more scalable, robust and realistic theory of socioeconomic systems.

Acknowledgements

Georgios Piliouras and Jeff S. Shamma would like to acknowledge AFOSR/MURI Project FA9550-09-1-0538, AFOSR Project FA9550-09-1-0420 as well as ONR Grant N00014-09-1-0751 for their support in the funding of this research. Georgios Piliouras would also like to acknowledge the CMI Wally Baer and Jeri Weiss postdoctoral fellowship, SUTD grant SRG ESD 2015 097 and MOE AcRF Tier 2 Grant 2016-T2-1-170, “A dynamical systems approach to algorithmic game theory”. Evdokia Nikolova would like to acknowledge NSF CCF Grants 1216103, 1350823, 1331863. Part of the work was completed while Georgios Piliouras was a CMI Wally Baer and Jeri Weiss postdoctoral fellow at California Institute of Technology. Part of the work was completed while Georgios Piliouras and Evdokia Nikolova were resident scientists at the Simons Institute for the Theory of Computing.

REFERENCES

- AGHASSI, M. AND BERTSIMAS, D. 2006. Robust game theory. *Math. Program.* 107, 1, 231–273.
- ALAND, S., DUMRAUF, D., GAIRING, M., MONIEN, B., AND SCHOPPMANN, F. 2006. *STACS 2006: 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23-25, 2006. Proceedings*. Springer Berlin Heidelberg, Berlin, Heidelberg, Chapter Exact Price of Anarchy for Polynomial Congestion Games, 218–229.
- ANGELIDAKIS, H., FOTAKIS, D., AND LIANEAS, T. 2013. Stochastic congestion games with risk-averse players. In *Algorithmic Game Theory*. Lecture Notes in Computer Science Series, vol. 8146. Springer Berlin Heidelberg, 86–97.
- ARTZNER, P., DELBAEN, F., EBER, J.-M., AND HEATH, D. 1999. Coherent measures of risk. *Math. Finance* 9, 203–228.
- 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. ACM, New York, NY, USA, 285–294.

- BALCAN, M.-F., CONSTANTIN, F., AND EHRLICH, S. 2011. The snowball effect of uncertainty in potential games. In *Internet and Network Economics*. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1–12.
- BEN-TAL, A., GHAOUI, L., AND NEMIROVSKI, A. 2009. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press.
- BERTSIMAS, D., BROWN, D. B., AND CARAMANIS, C. 2011. Theory and applications of robust optimization. *SIAM review* 53, 3, 464–501.
- BIRGE, J. R. AND LOUVEAUX, F. 1997. *Introduction to Stochastic Programming* Corrected Ed. Springer Series in Operations Research and Financial Engineering. Springer.
- BOT, R., LORENZ, N., AND WANKA, G. 2007. *Chance-constrained Optimization Problems and Applications*. Preprint. Techn. Univ., Fak. für Mathematik.
- CARAGIANNIS, I., FANELLI, A., GRAVIN, N., AND SKOPALIK, A. 2011. Efficient computation of approximate pure nash equilibria in congestion games. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. FOCS '11. IEEE Computer Society, Washington, DC, USA, 532–541.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND TARDOS, É. 2012. On the efficiency of equilibria in generalized second price auctions. *CoRR abs/1201.6429*.
- CHRISTODOULOU, G. AND KOUTSOPIAS, E. 2005. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. STOC '05. ACM, New York, NY, USA, 67–73.
- CHRISTODOULOU, G., KOVACS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *ICALP*. 820–832.
- CONWAY, J. H. AND GUY, R. 1996. *The Book of Numbers*. Springer.
- FIAT, A. AND PAPADIMITRIOU, C. 2010. When the players are not expectation maximizers. In *Proceedings of the Third international conference on Algorithmic game theory*. SAGT'10. Springer-Verlag, Berlin, Heidelberg, 1–14.
- FILAR, J. AND VRIEZE, K. 1996. *Competitive Markov Decision Processes*. Springer.
- GAIRING, M., MONIEN, B. AND TIEMANN, K. 2005. Selfish routing with incomplete information. In *Proceedings of the 17th Annual ACM Symposium on Parallel Algorithms and Architectures*. 203212.
- GARG, D. AND NARAHARI, Y. 2005. Price of anarchy of network routing games with incomplete information. In *Internet and Network Economics*, X. Deng and Y. Ye, Eds. Lecture Notes in Computer Science Series, vol. 3828. Springer Berlin Heidelberg, 1066–1075.
- GEORGIU, C., PAVLIDES, T., AND PHILIPPOU, A. 2006. Network uncertainty in selfish routing. In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*. 105–105.
- HARSANYI, J. 1967. Games with incomplete information played by bayesian players. *Management Science* 14, 8, 159–182, 320–334, 486–502.
- HEYDENREICH, B., MLLER, R., AND UETZ, M. 2007. Games and mechanism design in machine scheduling an introduction. *Production and Operations Management* 16, 4, 437–454.
- HOTA, A., GARG, S., AND SUNDARAM, S. 2013. Resource sharing games with failures and heterogeneous risk attitudes. In *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*. 535–542.
- IMMORLICA, N., KALAI, A. T., LUCIER, B., MOITRA, A., POSTLEWAITE, A., AND TENNENHOLTZ, M. 2011. Dueling algorithms. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*. STOC '11. ACM, New York, NY, USA, 215–224.
- KOUTSOPIAS, E. AND PAPADIMITRIOU, C. 1999. Worst-case equilibria. In *Sympo-*

- sium on Theoretical Aspects of Computer Science (STACS)*. Springer-Verlag, 404–413.
- LIANEAS, T. 2014. Congestion games: Stochastic extensions and techniques for reducing the price of anarchy. Ph.D. thesis, National Technical University of Athens. Dept. of Electrical and Computer Engineering.
- LIGETT, K. AND PILIOURAS, G. 2011. Beating the best Nash without regret. *SIGecom Exchanges* 10.
- LUCIER, B. AND BORODIN, A. 2010. Price of anarchy for greedy auctions. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 537–553.
- LUCIER, B. AND PAES LEME, R. 2011. Gsp auctions with correlated types. In *EC*. 71–80.
- MACHINA, M. J. 1987. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives* 1, 1, 121–154.
- MARKOWITZ, H. 1952. Portfolio selection. *The Journal of Finance* 7, 1, 77–91.
- MAVRONICOLAS, M. AND MONIEN, B. 2012. Minimizing expectation plus variance. In *Proceedings of the 5th international conference on Algorithmic Game Theory. SAGT'12*. Springer-Verlag, Berlin, Heidelberg, 239–250.
- MCKELVEY, R. D. AND PALFREY, T. R. 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10, 1, 6 – 38.
- MEIR, R. AND PARKES, D. 2015a. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *SIGMETRICS Perform. Eval. Rev.* 43, 3, 67–70.
- MEIR, R. AND PARKES, D. C. 2015b. Congestion games with distance-based strict uncertainty. In *AAAI*. AAAI Press, 969–974.
- MIRROKNI, V. S., THAIN, N., AND VETTA, A. 2012. A theoretical examination of practical game playing: Lookahead search. In *SAGT*, M. Serna, Ed. Lecture Notes in Computer Science Series, vol. 7615. Springer, 251–262.
- NASH, J. 1951. Non-cooperative games. *The Annals of Mathematics* 54, 2, 286–295.
- NEUMANN, J. V. AND MORGENSTERN, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- NEYMAN, A. AND SORIN, S. 2003. *Stochastic Games and Applications*. Nato Science Series C. Springer.
- NIKOLOVA, E. 2010. Approximation algorithms for reliable stochastic combinatorial optimization. In *Proceedings of the 13th International Conference on Approximation. APPROX/RANDOM'10*. Springer-Verlag, Berlin, Heidelberg, 338–351.
- NIKOLOVA, E., KELNER, J. A., BRAND, M., AND MITZENMACHER, M. 2006. Stochastic shortest paths via quasi-convex maximization. In *Lecture Notes in Computer Science 4168 (ESA 2006)*. Springer-Verlag, 552–563.
- NIKOLOVA, E. AND STIER-MOSES, N. E. 2011. Stochastic selfish routing. In *Proceedings of the 4th international conference on Algorithmic game theory. SAGT'11*. Springer-Verlag, Berlin, Heidelberg, 314–325.
- NIKOLOVA, E. AND STIER-MOSES, N. E. 2014a. The burden of risk aversion in mean-risk selfish routing. *CoRR abs/1411.0059*.
- NIKOLOVA, E. AND STIER-MOSES, N. E. 2014b. A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research* 62, 2, 366–382.
- PAES LEME, R. AND TARDOS, E. 2010. Pure and bayes-nash price of anarchy for generalized second price auction. In *FOCS*.
- PANAGEAS, I. AND PILIOURAS, G. 2016. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *ACM Conference on Economics and Computation*.
- PAPADIMITRIOU, C. AND PILIOURAS, G. 2016. From nash equilibria to chain recurrent

- sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. ITCS '16. ACM, New York, NY, USA, 227–235.
- PRÉKOPA, A. 1995. *Stochastic Programming*. Mathematics and Its Applications. Springer.
- PUTERMAN, M. L. 1994. *Markov Decision Processes*. Wiley.
- ROCKAFELLAR, T. 2007. Coherent approaches to risk in optimization under uncertainty. *Tutorials in operations research*, 38–61.
- ROSENTHAL, R. W. 1973. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2, 65–67.
- ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*. STOC '09. ACM, New York, NY, USA, 513–522.
- ROUGHGARDEN, T. 2012. The price of anarchy in games of incomplete information. In *ACM Conference on Electronic Commerce*, B. Faltings, K. Leyton-Brown, and P. Ipeirotis, Eds. ACM, 862–879.
- SAVAGE, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46, 55–67.
- SHAPIRO, A., DENTCHEVA, D., AND RUSZCZYŃSKI, A. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics.
- SHAPLEY, L. S. 1953. Stochastic games. *Proceedings of the National Academy of Sciences* 39, 10, 1095–1100.
- SRINIVASAN, A. 2007. Approximation algorithms for stochastic and risk-averse optimization. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Philadelphia, PA, USA, 1305–1313.
- SWAMY, C. 2011. Risk-averse stochastic optimization: probabilistically-constrained models and algorithms for black-box distributions. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '11. SIAM, 1627–1646.
- SYRGKANIS, V. 2012. Bayesian games and the smoothness framework. *CoRR abs/1203.5155*.
- URYASEV, S. 2000. *Probabilistic Constrained Optimization: Methodology and Applications*. Nonconvex Optimization and Its Applications. Springer.
- VON NEUMANN, J. AND MORGENSTERN, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- WALD, A. 1939. Contributions to the Theory of Statistical Estimation and Testing Hypotheses. *The Annals of Mathematical Statistics* 10, 299–326.

APPENDIX

A. MISSING PROOFS FROM SECTION 5

LEMMA A.1. *For every pair of nonnegative integers x, y , it holds*

$$(x + 2)y \leq \frac{1}{4}x(x + 1) + \frac{5}{4}y(y + 1).$$

Proof of Lemma A.1:

$$\begin{aligned} (x + 2)y &\leq \frac{1}{4}x(x + 1) + \frac{5}{4}y(y + 1) \Leftrightarrow \\ xy + 2y &\leq \frac{x^2}{4} + \frac{x}{4} + \frac{5}{4}y^2 + \frac{5}{4}y \Leftrightarrow \\ 0 &\leq \left(\frac{x}{2} - y\right)^2 + \frac{x}{4} + \frac{y}{4}(y - 3) \end{aligned} \quad (1)$$

The last inequality trivially holds for all nonnegative integers x, y where $y \neq 1, 2$. So, it suffices to check these two special cases separately.

For $y = 1$, inequality 1 reduces to

$$\begin{aligned} 0 &\leq \left(\frac{x}{2} - 1\right)^2 + \frac{x}{4} - \frac{1}{2} \Leftrightarrow \\ 0 &\leq x^2 - 3x + 2 \Leftrightarrow \\ 0 &\leq (x - 1)(x - 2) \end{aligned}$$

However, the last inequality is indeed true for all (nonnegative) integers x . For $y = 2$, inequality 1 reduces to

$$\begin{aligned} 0 &\leq \left(\frac{x}{2} - 2\right)^2 + \frac{x}{4} - \frac{1}{2} \Leftrightarrow \\ 0 &\leq x^2 - 7x + 14 \end{aligned}$$

However, the last inequality is true for all x . \square

B. PROOFS FOR SECTION 9: PRICE OF ANARCHY ANALYSIS FOR WALD'S MINIMAX PRINCIPLE

LEMMA B.1. *For $\alpha, \beta \in \mathbb{R}$:*

$$\beta(\alpha + 1) = \frac{1}{2}\alpha\frac{\alpha + 1}{2} + 2\beta\frac{\beta + 1}{2} - \frac{\alpha}{4} - \frac{(\alpha - 2\beta)^2}{4}.$$

Proof of Lemma B.1:

$$\begin{aligned} \beta(\alpha + 1) &= \frac{1}{2}\alpha\frac{\alpha + 1}{2} + 2\beta\frac{\beta + 1}{2} - \frac{\alpha}{4} - \frac{(\alpha - 2\beta)^2}{4} \Leftrightarrow \\ \beta(\alpha + 1) + \frac{\alpha}{4} + \frac{(\alpha - 2\beta)^2}{4} &= \frac{1}{4}\alpha(\alpha + 1) + \beta(\beta + 1) \Leftrightarrow \\ 4\beta(\alpha + 1) + \alpha + (\alpha - 2\beta)^2 &= \alpha(\alpha + 1) + 4\beta(\beta + 1) \Leftrightarrow \\ &(\alpha - 2\beta)^2 = (\alpha - 2\beta)^2 \end{aligned}$$

\square

LEMMA B.2. For nonnegative integers $x_e, x_e^* \geq 0$:

$$\frac{1}{2} \sum_e \left(\frac{3}{4} a_e + \frac{b_e}{2} \right) x_e - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 \leq \frac{1}{4} \left(\sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e b_e x_e \right).$$

Proof of Lemma B.2:

$$\begin{aligned} \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + \frac{b_e}{2} \right) x_e - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 &\leq \frac{1}{4} \left(\sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e b_e x_e \right) \Leftrightarrow \\ \frac{1}{2} \sum_e \frac{3}{4} a_e x_e - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 &\leq \frac{1}{4} \sum_e a_e \frac{x_e(x_e + 1)}{2} \Leftrightarrow \\ 3 \sum_e a_e x_e &\leq \sum_e a_e x_e (x_e + 1) + \sum_e a_e (x_e - 2x_e^*)^2 \end{aligned}$$

However, for all nonnegative integers $x_e \neq 1, 3x_e \leq x_e(x_e + 1)$, hence over those edges e the inequalities trivially hold. Even if $x_e = 1$, we have that $3x_e \leq x_e(x_e + 1) + (x_e - 2x_e^*)^2$, since this translates now to $(1 - 2x_e^*)^2 \geq 1$, which is true for all integers x_e^* . \square

Proof of Theorem 9.2: We will construct a congestion game for $N \geq 3$ players and $|E| = 2N$ facilities where the ratio between worst potential at equilibrium and its optimal value is equal to 2. All facilities e have latency function $c_e(x) = x$. We divide the set E into two subsets $E_1 = \{h_1, \dots, h_N\}$ and $E_2 = \{g_1, \dots, g_N\}$, each of N facilities. Player i has two pure strategies: $\{h_i, g_i\}$ and $\{h_{i+1}, g_{i-1}, g_{i+1}\}$. The optimal Nash is for each player to select the first strategy while the worst-case Nash equilibrium is for each player to select the second strategy. It is not hard to verify that at the worst Nash equilibrium the potential is equal to $4N$. At the optimal Nash the potential is equal to $2N$. Hence, their respective ratio is equal to 2 matching the upper bound for the class of linear congestion games.

For any $\epsilon > 0$ we can even construct symmetric linear congestion games for which this ratio is greater than $2 - \epsilon$. We construct such games as follows: We partition the facilities into sets $\{P_1, P_2, \dots, P_n\}$ of the same cardinality and make each P_i a pure strategy. Each P_i contains $N\alpha_1 + \binom{N}{2}\alpha_2$ facilities where α_1, α_2 are appropriate constants to be determined later. We furthermore partition each of the $N\alpha_1$ facilities of set P_i into N groups $\{R_1^{P_i}, \dots, R_N^{P_i}\}$, each of size α_1 . Similarly, we partition the $\binom{N}{2}\alpha_2$ facilities of set P_i into $\binom{N}{2}$ groups $Q = \{Q_{12}^{P_i}, \dots, Q_{ij}^{P_i}, \dots, Q_{N-1N}^{P_i}\}$ each of size α_2 .

The available strategies to each player i will include, besides all the P_i strategies, strategies (A_1, \dots, A_n) of the following form: Each strategy A_i contains out of each P_j , all the facilities in $R_i^{P_j}$ as well as all facilities in each of the $Q_{kl}^{P_j}$ where either $k = i$ or $l = i$. As a result, each strategy A_i includes $N\alpha_1 + N(N-1)\alpha_2$ elements.

The strategy profile (P_1, P_2, \dots, P_n) is trivially a Nash equilibrium of the game, since each P_i is an action with a minimal number of facilities (amongst all available actions) and the load of each facility is optimal (i.e. equal to 1).

We will choose α_1, α_2 , so that the strategy profile (A_1, A_2, \dots, A_n) is also a Nash equilibrium. Now, the cost for each player i is $c_i(A) = N(\alpha_1 + 2(N-1)\alpha_2)$. It suffices to select α_1, α_2 so that player i will not switch to P_j . (It is straightforward that player i will not switch to the Nash strategy of some other player k .) The cost after switching is as follows:

$$c_i(A_{-i}, P_j) = \alpha_1 + 2(N-1)\alpha_2 + 2(N-1)\alpha_1 + 3\binom{N-1}{2}a_2 = (2N-1)\alpha_1 + (N-1)\frac{3N-2}{2}\alpha_2$$

We want $c_i(A) = c_i(A_{-i}, P_j)$, or equivalently $\alpha_1 = \frac{N+2}{2}\alpha_2$, which is satisfied when we select $\alpha_1 = N+2$ and $\alpha_2 = 2$. With this, the potential at the (A_1, A_2, \dots, A_n) Nash equilibrium is $N(N\alpha_1 + 3 \times \frac{N(N-1)}{2}\alpha_2) = N^2(4N-1)$. The potential at the (P_1, P_2, \dots, P_n) Nash, is equal to the social welfare, equal to $N^2(2N+1)$. As N goes to infinity this ratio goes to 2. \square

C. PROOFS FOR SECTION 10: PRICE OF ANARCHY ANALYSIS FOR SAVAGE'S MINIMAX PRINCIPLE

Proof of Theorem 10.2: We denote by $\text{Social-Cost}(s) = \sum_e \sum_{i=1}^{x_e} (a_e i + b_e)$ the true social cost as is recorded by the social designer. When identifying (Nash) equilibria for this game we will be utilizing the formulation of the game given by latency functions $c_{e_0}(x) = c_{e_0}(x)$ and $c_e(x) = c_e(x) + \Delta$ for all $e' \neq e$. That is, for every agent i , $C_i(e_0, s_{-i}) = c_{e_0}(\ell_{e_0}(s))$ and for each machine e other than e_0 , $C_i(e, s_{-i}) = c_e(\ell_e(s)) + \Delta$.

Suppose that s_1, s_2, \dots, s_n is a Nash equilibrium of a symmetric linear load balancing game and $s_1^*, s_2^*, \dots, s_n^*$ corresponds to the social optimum, then we have that:

$$\begin{aligned} \text{Social-Cost}(s_1, s_2, \dots, s_n) &= \sum_e \sum_{i=1}^{x_e} (a_e i + b_e) = \sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e b_e x_e \\ &= \frac{1}{2} \sum_e (a_e x_e + b_e) x_e + \frac{1}{2} \sum_e (a_e + b_e) x_e \leq \frac{1}{2} \sum_i C_i(s_i^*, s_{-i}) + \frac{1}{2} \sum_e (a_e + b_e) x_e \\ &\leq \frac{1}{2} (a_{e_0}(x_{e_0} + 1) + b_{e_0}) x_{e_0}^* + \frac{1}{2} \sum_{e \neq e_0} (a_e(x_e + 1) + b_e + \Delta) x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e \\ &\leq \frac{1}{2} \sum_e (a_e(x_e + 1) + b_e) x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e \\ &\leq \frac{1}{2} \sum_e (a_e(x_e + 1) + b_e) x_e^* + \frac{1}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) + \frac{1}{2} \sum_e (a_e + b_e) x_e \\ &= \frac{1}{2} \sum_e a_e x_e^*(x_e + 1) + \frac{1}{2} \sum_e b_e x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e + \frac{1}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) \\ &\leq \frac{1}{2} \left(\frac{1}{2} \sum_e a_e \frac{x_e(x_e + 1)}{2} + 2 \sum_e a_e \frac{x_e^*(x_e^* + 1)}{2} - \frac{1}{4} \sum_e a_e x_e - \frac{1}{4} \sum_e a_e (x_e - 2x_e^*)^2 \right) \quad (2) \\ &+ \frac{1}{2} \sum_e b_e x_e^* + \frac{1}{2} \sum_e (a_e + b_e) x_e + \frac{1}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) \\ &= \frac{1}{4} \sum_e a_e \frac{x_e(x_e + 1)}{2} + \sum_e a_e \frac{x_e^*(x_e^* + 1)}{2} + \frac{1}{2} \sum_e b_e x_e^* - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 \\ &+ \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + b_e \right) x_e + \frac{1}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{4} \sum_e \left(a_e \frac{x_e(x_e + 1)}{2} + b_e x_e \right) + \sum_e \left(a_e \frac{x_e^*(x_e^* + 1)}{2} + \frac{1}{2} b_e x_e^* \right) - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 \\
&\quad + \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + \frac{b_e}{2} \right) x_e + \frac{1}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) \\
&\leq \frac{1}{4} \text{Social-Cost}(s_1, \dots, s_n) + \frac{3}{2} \text{Social-Cost}(s_1^*, \dots, s_n^*) - \frac{1}{8} \sum_e a_e (x_e - 2x_e^*)^2 + \frac{1}{2} \sum_e \left(\frac{3}{4} a_e + \frac{b_e}{2} \right) x_e
\end{aligned}$$

Above, line 2 is derived by applying lemma B.1. By applying Lemma B.2, we have that:

$$\text{Social-Cost}(s_1, s_2, \dots, s_n) \leq \frac{1}{2} \text{Social-Cost}(s_1, s_2, \dots, s_n) + \frac{3}{2} \text{Social-Cost}(s_1^*, s_2^*, \dots, s_n^*)$$

The truth of the theorem follows immediately. \square

Proof of Theorem 10.3: Consider a load balancing game with $2M - 1$ agents with two machines, the first with a linear latency function of $c_1(x) = x$ and the second with constant cost function $c_2(x) = M$. We will argue that under Savage's minimax principle, the outcome where all agents are assigned to the first machine is a Nash equilibrium. Indeed, the maximum possible regret for any agent in that configuration is $(2M - 1) - M = M - 1$. On the other hand, if an agent was to deviate to the second machine then his maximum possible regret would be equal to $M - 1$ corresponding to the scenario where the scheduler of the first machine would order him first. Hence, this is a Nash equilibrium with social cost of $M(2M - 1)$. On the other hand, the social cost of the outcome assigning exactly M agents to the first machine is equal to $M(M + 1)/2 + M(M - 1)$. The ratio of these costs converges to $4/3$ as M grows to infinity.

The construction is the same as in Theorem 5.5. We construct a congestion game with $N \geq 3$ players and $|E| = 2N$ facilities. We divide the set E into two subsets $E_1 = \{h_1, \dots, h_N\}$ and $E_2 = \{g_1, \dots, g_N\}$, each of N facilities. Player i has two pure strategies: $\{h_i, g_i\}$ and $\{h_{i+1}, g_{i-1}, g_{i+1}\}$. The latency functions of the elements in E_1 are $c(x) = 2x$, whereas the latency functions of the elements in E_2 are $c(x) = x$. The optimal outcome corresponds to the state where each player selects the first strategy while the worst-case Nash equilibrium is for each player to select the second strategy. The second outcome is a Nash, since the worst case regret for each agent in this configuration is $6 - 2 = 4$, whereas if he were to deviate, the worst case regret would now be $7 - 3 = 4$. At the worst Nash equilibrium the actual social cost is equal to $5N$. At the optimal state the actual social cost is equal to $3N$, which results to a ratio of $5/3$. \square

D. EXTENSIONS TO POLYNOMIAL COST FUNCTIONS

In this section, we present upper bounds on the price of anarchy in the case where the latency functions are polynomial functions of degree p with nonnegative coefficients. As with the main section, we will assume our standard model of randomizing schedulers where all schedulers process the users according to a common and uniformly randomly chosen permutation. Moreover, these results can still be straightforwardly generalized to other models of randomizing schedulers as discussed in section 4.2. Our results are aggregated in Table III. Unlike the case of linear congestion games our results here are not tight and stronger specialized techniques are needed to capture the precise price of anarchy in these settings. Nevertheless, this section indicates the robustness

Table III. Upper bounds on the price of anarchy of p -degree polynomial congestion games under uniform schedulers.

Principle	Price of Anarchy
Minimizing Expected cost ¹⁰	$p^{O(p)}$
Average case analysis	$p^{O(p)}$
Win-or-Go-Home	unbounded
Second moment method	unbounded
Wald's minimax cost	$p^{O(p)}$
Savage's minimax regret ¹¹	$p^{O(p)}$

of our main conclusion that efficiency depends critically on the risk attitudes of the agents. Furthermore, the risk attitudes that lead to near efficient systems for linear congestion games still work well with polynomial cost functions.

Risk neutral optimization

Under this assumption for polynomial cost functions $c_e(x) = \sum_{k=1}^p a_{e,k} x^k$ with $a_{e,k} \geq 0$, the implied cost of an agent i at outcome s is:

$$\begin{aligned} C_i(s) &= \mathbb{E}_{\omega \sim p} [c_i(s, \omega)] = \sum_{e \in s_i} \mathbb{E}_{\omega \sim p} \left[c_e \left(\sum_{j: e \in s_j \wedge r_e(j) \leq r_e(i)} 1 \right) \right] = \sum_{e \in s_i} \frac{\sum_{i=1}^{\ell_e} c_e(i)}{\ell_e(s)} \\ &= \sum_{e \in s_i} \sum_{k=1}^p a_{e,k} \frac{\sum_{i=1}^{\ell_e(s)} i^k}{\ell_e(s)} \end{aligned}$$

It is well known that $\frac{\sum_{i=1}^{\ell_e(s)} i^k}{\ell_e(s)}$ is a polynomial of degree p . This is derived by Faulhaber's formula [Conway and Guy 1996]. As we explain in section 4.4, $\text{Social-Cost}(s) = \sum_e \sum_{i=1}^{\ell_e} c_e(i)$, which in this case translates to $\sum_i C_i(s)$, the standard sum of the individual cost functions. So, this case translates to a standard congestion game with polynomial latency functions. Now, if all the $\frac{\sum_{i=1}^{\ell_e(s)} i^k}{\ell_e(s)}$ coefficients of the resulting polynomial are nonnegative then the price of anarchy bound follows by applying well known results [Caragiannis et al. 2011; Aland et al. 2006]. This assumption is indeed true if all polynomial latency functions are at most cubic degrees. However, this is not true for all functions, e.g. if $c_e(x) = x^4$ then $\frac{\sum_{i=1}^{\ell_e(s)} i^k}{\ell_e(s)} = \frac{1}{30} (6n^4 + 15n^3 + 10n^2 - 1)$, whereas if $c_e(x) = x^4 + \frac{1}{15}x$ then $\frac{\sum_{i=1}^{\ell_e(s)} i^k}{\ell_e(s)} = \frac{1}{30} (6n^4 + 15n^3 + 10n^2 + n)$. The structure of Faulhaber polynomials is a subject of combinatorial interest as it is closely connected to Bernoulli numbers and producing tight price of anarchy guarantees for all possible polynomials of degree p in this case would require exploiting some beneficial aspects of their structure.

¹⁰This bound holds under specific assumptions on the set of allowable polynomial functions.

¹¹This bound holds in the case of symmetric load balancing games and as long as there exist resources e, e' such that $c_e(1) = c_{e'}(1)$.

Average case analysis

Under this assumption for polynomial cost functions $c_e(x) = \sum_{k=1}^p a_{e,k} x^k$ with $a_{e,k} \geq 0$, the implied cost of an agent i at outcome s is $\mathcal{C}_i = \sum_{e \in s_i} c_e\left(\frac{x+1}{2}\right)$. This is equivalent to having the agents play a standard congestion game where the new latency function of resource e is $c'_e = c_e\left(\frac{x+1}{2}\right)$; this is a polynomial of degree p with nonnegative coefficients. The social cost is once again defined as $\sum_e \sum_{i=1}^{\ell_e} c_e(i)$ (see section 4.4), however, unlike the risk neutral optimization case this is no longer equal to the sum of the implied costs of the agents. It has been proven in [Caragiannis et al. 2011] that for any congestion game with polynomial latency function of degree p the ratio between the maximum and minimum potential at equilibrium is at most $p^{O(p)}$. The potential function under latency functions $c_e\left(\frac{x+1}{2}\right)$ is equal to $\sum_e \sum_{i=1}^{\ell_e} c_e\left(\frac{i+1}{2}\right)$. The ratio between $\sum_e \sum_{i=1}^{\ell_e} c_e\left(\frac{i+1}{2}\right)$ and $\sum_e \sum_{i=1}^{\ell_e} c_e(i)$ is at most $O(2^p)$, therefore, the price of anarchy defined according to $\sum_e \sum_{i=1}^{\ell_e} c_e(i)$ is still at most $2^p \cdot p^{O(p)} = p^{O(p)}$.

Win-or-go home principle

The negative lower bounds results for the case of linear latency functions immediately carry over to this (more expansive) class.

Second moment stochastic optimization

The negative lower bounds results for the case of linear latency functions immediately carry over to this (more expansive) class.

Wald's minimax cost

As in the case of linear latency functions, analyzing Wald's minimax cost under polynomial cost functions reduces to the question of analyzing the ratio of the potential at the (pure) Nash equilibrium with the highest potential and the pure Nash equilibrium with the lowest potential. This bound is at most $p^{O(p)}$ [Caragiannis et al. 2011].

Savage's minimax regret

As in the case of linear cost functions, if we assume symmetric load balancing games and furthermore that there exist resources e, e' such that $c_e(1) = c_{e'}(1)$ then a strategy outcome is an equilibrium under Savage's minimax regret principle if and only if it is an equilibrium under Wald's minimax cost principle. In both cases, the social cost is accounted in the same way (see section 4.4). As a result, we derive an upper bound of $p^{O(p)}$ for the price of anarchy. Removing the structural assumption that $c_e(1) = c_{e'}(1)$ is non-trivial even under the assumption of linear cost functions and it is likely that more technical arguments will be needed to extend this to polynomial functions. Removing the others assumptions (*i.e.*, symmetry, load balancing games) and finding the exact price of anarchy remains an open question even for the linear case.

The results in this section, unlike the case of linear cost functions are not tight. The development of tight price of anarchy results even in the standard case of potential games has required the development of carefully crafted ad-hoc techniques along with families of tight examples (e.g. [Aland et al. 2006]). A complete understanding of the price of anarchy of polynomial congestion games under different risk attitudes will similarly require carefully designed analytical techniques and tools.