

# Congestion Games with Distance-Based Strict Uncertainty

Reshef Meir and David Parkes\*

Harvard University

## Abstract

We put forward a new model of congestion games where agents have uncertainty over the routes used by other agents. We take a non-probabilistic approach, assuming that each agent knows that the number of agents using an edge is within a certain range. Given this uncertainty, we model agents who either minimize their worst-case cost (WCC) or their worst-case regret (WCR), and study implications on equilibrium existence, convergence through adaptive play, and efficiency. Under the WCC behavior the game reduces to a modified congestion game, and welfare improves when agents have moderate uncertainty. Under WCR behavior the game is not, in general, a congestion game, but we show convergence and efficiency bounds for a simple class of games.

## Introduction

Congestion games (Rosenthal 1973) provide a good abstraction for a wide spectrum of scenarios where self-interested agents contest for resources, and can be conveniently analyzed using game-theoretic tools.

Recently, more complex models of congestion games have been suggested, taking into account the incomplete information agents may have when making a decision (e.g. Ashlagi et al. (2009) and Piliouras et al. (2013); see Related Work). Uncertainty may stem from multiple sources, including uncertainty about the state of nature— and thus the cost of resources —or about other agents’ actions. We can imagine commuters choosing routes home from work and facing uncertainty about road conditions (e.g., weather, roadworks) as well as about the routes selected by others.

Rather than model uncertainty through a distributional model, we adopt a non-probabilistic approach of *strict uncertainty*. Indeed, extensive experimental and empirical studies have demonstrated that people have difficulties in representing probabilities, and often adopt other heuristics in place of probabilistic reasoning (Tversky and Kahneman 1974; Slovic, Fischhoff, and Lichtenstein 1980). With strict uncertainty, each player faces a set of *possible states*, where the cost of each action depends on the (unknown) actual state, and decisions take into account risk attitudes

and other biases, often applying heuristics rather than optimization. Such alternative approaches to decision making in general, and to uncertainty in particular, have deep roots in the AI literature, largely due to the works of Herbert Simon on *bounded rationality* and *procedural rationality* (Simon 1957; 1987).

Having adopted a non-probabilistic approach, we must make two crucial modeling decisions. First, we must decide how each agent acts in the face of strict uncertainty. The simplest behavior follows a *minimax approach* (Wald 1939; Simon 1957), and assumes that the decision maker is trying to minimize her *worst-case cost* (WCC). Another approach seeks to minimize the *worst-case regret* (WCR) of the decision maker, which goes back to Savage (1951), and has been also applied to games (Hyafil and Boutilier 2004). Both cost measures are worst-case approaches, and suitable as an abstraction for the behavior of a rational but risk-averse agent.

Second, we need to determine which states are considered possible by the agents. To construct the set of possible states, we adopt the recent model of *distance-based uncertainty* (Meir, Lev, and Rosenschein 2014). All agents share the same belief about the current “reference state” of the network (i.e., the load on every edge in a routing game), which may be available for example from an external source such as traffic reports, or from an agent’s previous experience. However agents vary in the accuracy they attribute to the reference state. Each agent  $i$  has an uncertainty parameter,  $r_i$ , which reflects a belief that the actual load is within some distance  $r_i$  of the reference load. A higher  $r_i$  may reflect either that an agent is less informed about the true congestion, or, alternatively, that she is more risk-averse.

From each heuristic (WCC or WCR) we can derive a natural equilibrium concept. Intuitively, every action profile induces a reference state  $s$ , and we consider the heuristic best response of every agent to the set of possible states around  $s$ . State  $s$  is an *equilibrium* if every agent minimizes her worst case cost (or regret) by keeping her current action.

As a simple example, if in some profile 100 players are using a resource, then agent  $i$  believes the actual load to be anywhere between  $100/r_i$  and  $100r_i$ . If  $r_i = 1$  for all agents, we get the standard complete information model as a special case: both minimax cost and minimax regret collapse to simple cost minimization, and our equilibrium notion coincides with Nash equilibrium.

\*Parkes is supported in this work by the TomKat Charitable Trust.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

**Our contribution.** We study equilibrium behavior in *nonatomic congestion games*, under our strict, distance-based model of uncertainty. For simplicity and concreteness, we focus our presentation on routing games, where resources are edges in a graph, and valid strategies are paths from source to target.<sup>1</sup> With worst-case cost players, we show that the game reduces to a modified, complete-information routing game with *player-specific costs*. Further, if all agents have the same uncertainty level we get a potential game.

We are interested in how equilibrium quality (measured by the price of anarchy (Roughgarden and Tardos 2004)) is affected by introducing uncertainty. For routing games with affine cost functions, we show that the price of anarchy (PoA) under uncertainty decreases gradually from  $\frac{4}{3}$  (without uncertainty) to 1, and then climbs back up, proportionally to the amount of uncertainty. We also show that in a population of agents with different uncertainty levels, the PoA is bounded by the PoA of the worst possible uncertainty level.

With worst-case regret players the induced game is no longer a congestion game. Yet, we show that for a simple class of games a weak potential function exists, and thus equilibrium existence and convergence results are available. We give some preliminary results on PoA bounds with worst-case regret players. Due to space constraints most of our proofs are omitted, and are available in the full version of this paper.<sup>2</sup>

## Preliminaries

**Nonatomic routing games.** Following Roughgarden (2003) and Roughgarden and Tardos (2004), a *nonatomic routing game* (NRG) is a tuple  $\mathcal{G} = \langle G, \mathbf{c}, m, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$ , where

- $G = (V, E)$  is a directed graph;
- $\mathbf{c} = (c_e)_{e \in E}$ ,  $c_e(t) \geq 0$  is the cost incurred when  $t$  agents use edge  $e$ ;
- $m \in \mathbb{N}$  is the number of agent types;
- $\mathbf{u}, \mathbf{v} \in V^m$ , where  $(u_i, v_i)$  are the source and target nodes of type  $i$  agents;
- $\mathbf{n} \in \mathbb{R}_+^m$ , where  $n_i \in \mathbb{R}_+$  is the total mass of type  $i$  agents.  $n = \sum_{i \leq m} n_i$  is the total mass of agents.

We denote by  $A_i \subseteq 2^E$  the set of all directed paths between the pair of nodes  $(u_i, v_i)$  in the graph. Thus  $A_i$  is the set of *actions* available to agents of type  $i$ . We denote by  $A = \cup_i A_i$  the set of all directed source-target paths. We assume that the costs  $c_e$  are non-decreasing, continuous and differentiable.

A NRG is *symmetric* if all agents have the same source and target, i.e.,  $A_i = A$  for all  $i$ . A symmetric NRG is a *resource selection game* (RSG) if  $G$  is a graph of parallel links. That is, if  $A = E$  and the action of every agent is to select a single resource (edge).

**Game states.** A *state* (or action profile) is a vector  $\mathbf{s} \in \mathbb{R}_+^{|A| \times m}$ , where  $s_{f,i}$  is the amount of agents of type  $i$  that use path  $f \in A_i$ . In a valid state,  $\sum_{f \in A_i} s_{f,i} = n_i$ . The total traffic on path  $f \in A$  is denoted by  $s_f = \sum_{i=1}^m s_{f,i}$ . The *load state*  $\bar{\mathbf{s}} \in \mathbb{R}_+^{|E|}$  is a vector of aggregated edge loads derived from state  $\mathbf{s}$ , where  $\bar{s}_e = \sum_{f: e \in f} s_f$ . This is the total traffic on edge  $e \in E$  via all paths going through  $e$ . We also allow load states that *cannot be derived* from a valid state.

Note that all the information relevant to the costs in state  $\mathbf{s}$  is specified in the load state  $\bar{\mathbf{s}}$ : all agents using a particular edge  $e$  suffer a cost of  $c_e(\bar{s}_e)$  in state  $\mathbf{s}$ , and the cost of using a path  $f \in A_i$  is  $c(f, \bar{\mathbf{s}}) = \sum_{e \in f} c_e(\bar{s}_e)$ . Thus except in settings where agents' types or exact strategies matter, we may use  $\mathbf{s}$  and  $\bar{\mathbf{s}}$  interchangeably. The *social cost* in a profile  $\mathbf{s}$  is

$$SC(\mathbf{s}) = \sum_{i=1}^m \sum_{f \in A_i} s_{f,i} c(f, \bar{\mathbf{s}}) = \sum_{e \in E} \bar{s}_e c_e(\bar{s}_e).$$

The second equality is since we multiply the cost  $c(t)$  by the mass  $t$  ("number of agents") who experience it.

**Equilibrium and potential.** Without uncertainty, a state  $\mathbf{s}$  for an NRG is an *equilibrium* if for every agent type  $i$  and actions  $f_1, f_2 \in A_i$  with  $s_{i,f_1} > 0$ ,  $c(f_1, \bar{\mathbf{s}}) \leq c(f_2, \bar{\mathbf{s}})$ . That is, if no agent can switch to a path with a lower cost. This is the analogy of a Nash equilibrium in nonatomic games.

In nonatomic games,  $\phi(\mathbf{s})$  is a *potential function*, if any (infinitesimally small) rational move, i.e., a move that decreases the cost of the moving agents, also lowers the potential.  $\phi(\mathbf{s})$  is a *weak potential function* if at any state there is at least one such move (although some rational moves may increase  $\phi$ ). Any game with a potential function is *acyclic*, in the sense that such "infinitesimal best responses" of self interested agents are guaranteed to converge to a local minimum of the potential function (and an equilibrium). A game with a weak potential may have cycles, but from any state there is some path of rational moves that leads to an equilibrium.

It is well known that NRGs have a potential function, which is defined as (we omit the argument  $\mathcal{G}$  when it is clear from the context):  $\phi(\mathbf{s}) = \phi(\mathcal{G}, \mathbf{s}) = \sum_{e \in E} \int_{t=0}^{\bar{s}_e} c_e(t) dt$ . Furthermore, in a NRG every local minimum of the potential is also a global minimum; all equilibria have the same social cost; and in every equilibrium all agents of type  $i$  experience the same cost (Aashtiani and Magnanti 1981; Milchtaich 2000; Roughgarden and Tardos 2004).

**Affine routing games.** In an *affine* NRG, all cost functions take the form of a linear function. That is,  $c_e(t) = a_e t + b_e$  for some constants  $a_e, b_e \geq 0$ . In an affine game  $\mathcal{G}$ , the social cost can be written as  $SC(\mathcal{G}, \mathbf{s}) = \sum_{e \in E} a_e (\bar{s}_e)^2 + b_e \bar{s}_e$ ; and the potential as  $\phi(\mathcal{G}, \mathbf{s}) = \sum_{e \in E} \frac{1}{2} a_e (\bar{s}_e)^2 + b_e \bar{s}_e$ . *Pigou's example* is the special case of an affine RSG with two resources, where  $c_1(t) = 1$  and  $c_2(t)$  is defined with  $b_2 = 0$ . We will use variations of this example throughout the paper, and denote by  $\mathcal{G}_P(a_2, n)$  the instance where  $c_2(t) = a_2 t$ , and there is a mass of  $n$  agents.

<sup>1</sup>Any congestion game is equivalent to a routing game.

<sup>2</sup><http://arxiv.org/abs/1411.4943>

**Potential and social cost.** The social cost of every NRG can be written as the potential of a suitably modified game. For this, let  $\hat{\mathcal{G}}$  be a modification of  $\mathcal{G}$ , where we replace every  $c_e(t)$  with  $\hat{c}_e(t) = c_e(t) + tc'_e(t)$ . Then,  $\phi(\hat{\mathcal{G}}, \mathbf{s}) = SC(\mathcal{G}, \mathbf{s})$  for all  $\mathbf{s}$  (Roughgarden 2007). For an affine game, the modified cost function is  $\hat{c}_e(t) = 2a_e t + b_e$ ; and  $\phi(\hat{\mathcal{G}}, \mathbf{s}) = \sum_{e \in E} a_e (\bar{s}_e)^2 + b_e \bar{s}_e = SC(\mathcal{G}, \mathbf{s})$ .

**The price of anarchy.** Let  $EQ(\mathcal{G})$  be the set of equilibria in game  $\mathcal{G}$ . The *price of anarchy* (PoA) of a game is the ratio between the social cost in the worst equilibrium in  $EQ(\mathcal{G})$  and the optimal social cost. Since all equilibria have the same cost, we can write  $\text{PoA}(\mathcal{G}) = \frac{SC(\mathbf{s}^*)}{SC(OPT)}$ , where  $\mathbf{s}^*$  is an arbitrary equilibrium of  $\mathcal{G}$ . In affine NRGs, it is known that  $\text{PoA}(\mathcal{G}) \leq \frac{4}{3}$ , and this bound is attained by  $\mathcal{G}_P(1, 1)$  (Roughgarden and Tardos 2004).

### Introducing uncertainty

In our strict uncertainty model, there is an underlying base game  $\mathcal{G}$ , which is a NRG. However given an action profile (state)  $\mathbf{s}$ , each agent believes that there is some set of possible states, and selects her action based on worst-case assumptions.

To define this set of possible states, we augment the description of every agent type with an *uncertainty parameter*  $r_i \geq 1$ , and denote  $\mathbf{r} = (r_i)_{i \leq m}$ . The special case where all  $r_i = r$  is called *homogeneous uncertainty*. We adopt distance-based uncertainty, so that in a given state  $\mathbf{s}$  (where the actual load on edge  $e$  is  $\bar{s}_e$ ), a type  $i$  agent believes that the load is anywhere in the range  $[\bar{s}_e/r_i, \bar{s}_e \cdot r_i]$ . Consequently, the agent believes that the cost she will suffer from using resource  $e$  is between  $c_e(\bar{s}_e/r_i)$  and  $c_e(\bar{s}_e \cdot r_i)$ . Agents apply this reasoning separately to each resource, thus the load state  $\bar{\mathbf{s}}'$  is considered *possible* in load state  $\bar{\mathbf{s}}$  by a type  $i$  agent, if  $\bar{s}'_e \in [\bar{s}_e/r_i, \bar{s}_e \cdot r_i]$  for all  $e \in E$ .<sup>3</sup>

In other words, consider the distance metric  $d(\bar{\mathbf{s}}, \bar{\mathbf{s}}') = \min\{x \geq 0 : \forall e \in E, \bar{s}_e \geq \frac{\bar{s}'_e}{1+x} \wedge \bar{s}'_e \geq \frac{\bar{s}_e}{1+x}\}$ . Then  $S(\mathbf{s}, r_i) = S(\bar{\mathbf{s}}, r_i) = \{\bar{\mathbf{s}}' \in \mathbb{R}_+^{|E|} : d(\bar{\mathbf{s}}, \bar{\mathbf{s}}') \leq r_i - 1\}$  is the set of load states that a type  $i$  agent believes possible given  $\bar{\mathbf{s}}$ . Note that  $\bar{\mathbf{s}}'$  may not correspond to any actual state  $\mathbf{s}'$ , e.g. the total load on all paths may not sum up to total mass  $n$ , as an agent may not know exactly how many other agents participate.

### Behavior and equilibria

**Worst-case cost.** Under the WCC model, each agent cares about the worst possible cost of each action. Thus for an agent of type  $i$ , the effective cost of choosing path (action)  $f \in A_i$  in state  $\mathbf{s}$  is  $c_i^*(f, \bar{\mathbf{s}}) = \max\{c(f, \bar{\mathbf{s}}') : \bar{\mathbf{s}}' \in S(\bar{\mathbf{s}}, r_i)\}$ .

<sup>3</sup>In the language of modal logic, we say that the  $\bar{\mathbf{s}}'$  is *accessible* from  $\bar{\mathbf{s}}$  if the above holds. Our accessibility relation is symmetric, but non-transitive. While transitivity is a well accepted axiom in epistemic models (Aumann 1999), we argue that it does not make sense when there is a natural metric over states.




Figure 1: Two resources with base costs  $c_1(\bar{s}_1) = 5 + \bar{s}_1, c_2(\bar{s}_2) = 2 + \bar{s}_2$ . The figure shows the true costs for  $\bar{s}_1 = 3, \bar{s}_2 = 5$ . The dotted brackets show the range of possible costs for  $r = 2$ , where the upper bracket is the WCC cost. The dashed lines are the WCR costs. We can see that the better resource under WCC is  $e_1$ , but  $e_2$  is better under WCR, as  $c^{**}(e_2, \bar{\mathbf{s}}) = 12 - 6.5 = 11 - 4.5 = c^{**}(e_1, \bar{\mathbf{s}})$ .

Every NRG  $\mathcal{G}$  and uncertainty vector  $\mathbf{r}$  induce a new nonatomic game  $\mathcal{G}^*(\mathbf{r})$ , where the cost functions are  $c_i^*$ . That is, a type  $i$  agent playing so as to minimize her worst-case cost in  $\mathcal{G}$ , behaves exactly like a “rational” type  $i$  agent (minimizing exact cost) in  $\mathcal{G}^*(\mathbf{r})$ . *A priori*,  $\mathcal{G}^*(\mathbf{r})$  is not a NRG, but it has a very similar structure. For every action  $f \in A_i$ , according to the WCC model,

$$c_i^*(f, \bar{\mathbf{s}}) = \sum_{e \in f} \max\{c_e(\bar{s}'_e) : \bar{s}'_e \in S(\bar{\mathbf{s}}, r_i)\} = \sum_{e \in f} c_e(r_i \bar{s}_e).$$

Since  $c_e(r_i t)$  can be written as a player-specific cost function,  $c_{i,e}^*(t)$ , we have that  $\mathcal{G}^*(\mathbf{r})$  is a NRG with *player-specific* costs (Milchtaich 2005), where each player type can adopt a different cost function; e.g., for affine games  $c_{i,e}^*(t) = r_i a_e t + b_e$ .

**Worst-case regret.** We get a different modified game,  $\mathcal{G}^{**}(\mathbf{r})$ , under the WCR model. The *regret* (for a type  $i$  agent) of playing action  $f$  in state  $\bar{\mathbf{s}}'$  is defined as  $REG_i(f, \bar{\mathbf{s}}') = c(f, \bar{\mathbf{s}}') - \min_{f' \in A_i} c(f', \bar{\mathbf{s}}')$ . Given this, the cost  $c_i^{**}(f, \bar{\mathbf{s}})$  in the modified game, which is the *worst-case regret* a type  $i$  agent may suffer for playing  $f$ , is defined as:

$$c_i^{**}(f, \bar{\mathbf{s}}) = \max\{REG_i(f, \bar{\mathbf{s}}') : \bar{\mathbf{s}}' \in S(\bar{\mathbf{s}}, r_i)\}.$$

This cost function  $c_i^{**}(f, \bar{\mathbf{s}})$  does not have a natural decomposition to edge-wise costs, since regret depends also on the load on unused edges. An example of WCC and WCR costs in a simple 2-resource RSG appear in Figure 1.

**Equilibrium.** A *WCC equilibrium* is a state where no agent can improve her worst-case cost w.r.t. her uncertainty level. By definition of the cost function  $c^*$ , the WCC equilibria of  $\mathcal{G}$  for uncertainty values  $\mathbf{r}$  are exactly the Nash equilibria of  $\mathcal{G}^*(\mathbf{r})$ . Similarly, a *WCR equilibrium* is a Nash equilibrium of  $\mathcal{G}^{**}(\mathbf{r})$ . Since both of  $\mathcal{G}^*(\mathbf{r})$  and  $\mathcal{G}^{**}(\mathbf{r})$  are special cases of nonatomic games, existence of equilibria follows from general existence theorems (Schmeidler 1973). However the other properties of NRG, such as the existence of a potential function, and bounds on the PoA, are not guaranteed.

## Routing Games with WCC players

**Equilibrium and convergence.** For the special case of  $r_i = r$  for all  $i \leq m$ , it is not hard to see that  $\mathcal{G}^*(r)$  is a non player-specific NRG. This is since  $c_{e,i}^*(\bar{\mathbf{s}}) = c_e(r \cdot \bar{s}_e)$  is only a function of  $\bar{s}_e$ . We denote this modified cost function by  $c_e^r(t)$ . It follows that  $\mathcal{G}^*(r)$  is a potential game, where  $\phi(\mathcal{G}^*(r), \mathbf{s}) = \sum_e \int_{t=0}^{\bar{s}_e} c_e^r(t) dt$ . Thus  $\mathcal{G}^*(r)$  is acyclic, the equilibria of  $\mathcal{G}^*(r)$  are the minima of  $\phi(\mathcal{G}^*(r), \mathbf{s})$ , and all equilibria have the same social cost.

The more interesting question is what properties of NRG are maintained when agents have different uncertainty parameters. We have already noted that in  $\mathcal{G}^*(\mathbf{r})$  there is at least one equilibrium. Player-specific RSGs are known to have a weak potential (Milchtaich 1996), but this does not preclude cycles. Indeed, we show that a cycle may occur even in an RSG where agents only differ in their uncertainty level.

**Proposition 1.** *There is an RSG  $\mathcal{G}$  with 3 resources, and a vector  $\mathbf{r}$  s.t.  $\mathcal{G}^*(\mathbf{r})$  contains a cycle.*

## Equilibrium quality for affine games

Recall that under the WCC model, agents play as if they take part of the game  $\mathcal{G}^*(\mathbf{r})$ , while their actual, realized costs are those in underlying game  $\mathcal{G}$ . We thus define the *Price of Anarchy for WCC players with uncertainty vector  $\mathbf{r}$*  as:

$$\text{C-PoA}(\mathcal{G}, \mathbf{r}) = \max_{\mathbf{s} \in \text{EQ}(\mathcal{G}^*(\mathbf{r}))} \frac{SC(\mathcal{G}, \mathbf{s})}{SC(\mathcal{G}, \text{OPT}(\mathcal{G}))}.$$

We focus our analysis on games with affine costs, and look for bounds on  $\text{C-PoA}(\mathcal{G}, \mathbf{r})$ . In particular, we explore whether players with uncertainty reach better or worse social outcomes under WCC behavior than under standard, complete information equilibria.

**Homogeneous uncertainty.** We start with the simplifying assumption that  $r_i = r$  for all types. Recall that in this case  $\mathcal{G}^*(r)$  is a non player specific NRG, where the cost of each edge is modified to  $c_e^r(t) = ra_e t + b_e$ . These modified costs can be attained in other contexts that do not involve uncertainty. For example, this can be achieved through taxation (Cole, Dodis, and Roughgarden 2003). Following the discussion in Potential and Social Cost, an optimal taxation scheme would perturb the cost functions so that the realized cost is  $\hat{c}_e(t) = c_e(t) + t c'_e(t)$ , as this will guarantee that

$\phi(\hat{\mathcal{G}}, \mathbf{s}) = SC(\mathcal{G}, \mathbf{s})$  for all  $\mathbf{s}$ . In this way, minimizing the potential of  $\hat{\mathcal{G}}$ , as happens in equilibrium, also minimizes the social cost in  $\mathcal{G}$  (see Section 18.3.1 in Roughgarden (2007) for a detailed explanation).

For the special case of affine games, it is easy to see that the effect of uncertainty level  $r = 2$  is equivalent to that of an optimal taxation scheme. That is,  $\phi(\mathcal{G}^*(2), \mathbf{s}) = SC(\mathcal{G}, \mathbf{s})$  for all  $\mathbf{s}$ . This means that if all agents adopt the WCC viewpoint for an uncertainty type of  $r = 2$ , then they would play the social optimum. Unfortunately, the value of  $r$  is not a design parameter—we cannot decide for the agents how uncertain they should be, since this reflects their beliefs. We would therefore like to have guarantees on equilibrium quality for any value of  $r$ . The next lemma provides our first result in this direction.

We denote  $\phi^r(\mathbf{s}) = \phi(\mathcal{G}^*(r), \mathbf{s})$ . It will be convenient to treat the cases  $r \geq 2$  and  $r \leq 2$  separately.

**Lemma 2.** *Let  $\mathcal{G}$  be an affine NRG, and suppose that  $r \geq 2$ . Then for all  $\bar{\mathbf{s}}$ ,  $\phi^r(\bar{\mathbf{s}}) \in [SC(\mathcal{G}, \bar{\mathbf{s}}), \frac{r}{2} SC(\mathcal{G}, \bar{\mathbf{s}})]$ .*

**Proposition 3.** *For  $r \geq 2$ , and any affine NRG  $\mathcal{G}$ ,  $\text{C-PoA}(\mathcal{G}, r) \leq \frac{r}{2}$ .*

*Proof.* Let  $\mathbf{s}^*$  be a global optimum of  $\phi^r$ ,  $\mathbf{s}_O = \text{OPT}(\mathcal{G})$ . By Lemma 2  $SC(\mathcal{G}, \mathbf{s}^*) \leq \phi^r(\mathbf{s}^*) \leq \phi^r(\mathbf{s}_O) \leq \frac{r}{2} SC(\mathcal{G}, \mathbf{s}_O)$ .  $\square$

We can similarly derive a bound of  $2/r$  for the range  $r \in [1, 2)$ , but we can do better. We next show that as we increase the uncertainty level  $r$  from 1 towards 2, we get a smooth improvement in social cost.

**Theorem 4.** *For  $r \in [1, 2]$ , and for any affine NRG  $\mathcal{G}$ ,  $\text{C-PoA}(\mathcal{G}, r) \leq 2 - \frac{2}{r} + (\frac{2}{r} - 1) \text{PoA}(\mathcal{G}) \leq \frac{2+2r}{3r}$ .*

Let  $q^*(r) = \max_{\mathcal{G}} \text{C-PoA}(\mathcal{G}, r)$ ; and  $q_P^*(r) = \max_{\mathcal{G}_P} \text{C-PoA}(\mathcal{G}_P, r)$ . The results above give us an upper bound on  $q^*(r)$ . By a careful analysis of the Pigou-type instances, we can compute  $q_P^*(r)$  exactly, which also provides us with a lower bound on  $q^*(r)$ .

**Corollary 5.** *For  $r \in [1, 2]$ ,  $q^*(r) \in [\frac{4}{4r-r^2}, \frac{2+2r}{3r}]$ . For  $r > 2$ , we have  $q^*(r) \in [\frac{r^2}{4(r-1)}, r/2] = [r/4 + o(1), r/2]$ . Also,  $q_P^*(r)$  is equal to our lower bound on  $q^*(r)$ .<sup>4</sup>*

Note that for  $r = 1$  and  $r = 2$ , we have  $q^*(r) = q_P^*(r)$ , and both are equal to the familiar values of  $4/3$  and 1, respectively. See Figure 2 for a graphical comparison of the bounds.

**Diverse population.** We would like to show that if we have a population of agents with different uncertainty levels, the social cost does not exceed that of the upper bound we have on the worst type; i.e., that  $q^*(\mathbf{r}) \leq q^*(r_i)$  for some type  $i$  in the mixture. We show something very close (our

<sup>4</sup>Roughgarden (2003) showed that for any class of cost functions, a worst-case example for the PoA can be constructed on an RSG with two resources. It is not clear if this is also true for the C-PoA, as in our case the optimum and the equilibrium are computed for two different games. However if a similar result can be proved, then our upper bounds would collapse to the lower bounds.




Figure 2: The solid red lines are upper- and lower-bounds on  $q^*(r)$ , i.e. the maximum PoA, across all games, for WCC players with uncertainty parameter  $r$  (the lower bound is also exactly  $q_P^*(r)$ , the maximum PoA across Pigou examples). The blue dashed line is exactly  $q_P^{**}(r)$ . The dotted lines mark  $4/3$  and  $1$ .

bound is slightly worse when there are types both below and above  $r = 2$ ).

Let  $j = \operatorname{argmin}_i r_i, k = \operatorname{argmax}_i r_i$ ; let  $\alpha_j = \text{C-PoA}(\mathcal{G}, r_j)$  if  $r_j < 2$  and  $1$  otherwise; let  $\alpha_k = \frac{r_k}{2}$  if  $r_k > 2$  and  $1$  otherwise.

**Lemma 6.** *Let  $\mathbf{s}, \mathbf{s}'$  be any two states in affine NRG  $\mathcal{G}$ , and consider  $r_3 \geq r_2 \geq r_1$ . If  $\phi^{r_3}(\mathbf{s}) \geq \phi^{r_3}(\mathbf{s}')$  and  $\phi^{r_2}(\mathbf{s}) \leq \phi^{r_2}(\mathbf{s}')$ , then  $\phi^{r_1}(\mathbf{s}) \leq \phi^{r_1}(\mathbf{s}')$  as well.*

**Theorem 7.** *Let  $\mathcal{G}$  be an affine NRG,  $\mathbf{r}$  be an uncertainty vector. Then  $\text{C-PoA}(\mathcal{G}, \mathbf{r}) \leq \alpha_j \cdot \alpha_k$ .*

*Proof of Theorem 7 for  $r_j \geq 2$ .* Let  $\mathbf{s}^*$  be an equilibrium of  $\mathcal{G}^*(\mathbf{r})$ . Let  $\mathbf{s}_i^*$  be a state minimizing  $\phi^{r_i}(\mathbf{s})$ . Note that  $\phi^r(\mathbf{s}) < \phi^{r'}(\mathbf{s})$  for all  $r < r'$  and any  $\mathbf{s}$ . By Lemma 2,  $\phi^2(\mathbf{s}) = SC(\mathcal{G}, \mathbf{s})$  for any  $\mathbf{s}$ . We next bound  $SC(\mathcal{G}, \mathbf{s}^*)$ , dividing into cases:  $r_j \geq 2; r_k \leq 2; r_j < 2 < r_k$ . We prove for  $r_j \geq 2$ .

Consider the state  $\mathbf{s}_k^*$ , which is an equilibrium of the game  $\mathcal{G}^*(r_k)$ . If  $\phi^{r_k}(\mathbf{s}^*) = \phi^{r_k}(\mathbf{s}_k^*)$ , then  $\mathbf{s}^*$  is also an equilibrium of  $\mathcal{G}^*(r_k)$ . Since  $\mathcal{G}^*(r_k)$  is an NRG, all equilibria have the same social cost, thus

$$\phi^{2r_k}(\mathbf{s}^*) = SC(\mathcal{G}^*(r_k), \mathbf{s}^*) = SC(\mathcal{G}^*(r_k), \mathbf{s}_k^*) = \phi^{2r_k}(\mathbf{s}_k^*).$$

As  $2r_k > r_k \geq 2$ , by Lemma 6  $\phi^2(\mathbf{s}_k^*) \geq \phi^2(\mathbf{s}^*)$  (in fact equal). Thus

$$SC(\mathbf{s}^*) = \phi^2(\mathbf{s}^*) \leq \phi^2(\mathbf{s}_k^*) = SC(\mathbf{s}_k^*). \quad (1)$$

Thus suppose that  $\phi^{r_k}(\mathbf{s}_k^*), \phi^{r_k}(\mathbf{s}^*)$  differ. By definition,  $\phi^{r_k}(\mathbf{s}_k^*) < \phi^{r_k}(\mathbf{s}^*)$ . There must be some (non zero measure of) agents with different actions in both states. It cannot be that all of these agents have uncertainty  $r_k$ , since the states  $\mathbf{s}^*, \mathbf{s}_k^*$  have a different  $\phi^k$  potential.

Consider any such agent of type  $i, r_i < r_k$ , whose action under  $\mathbf{s}^*$  differs from the one in  $\mathbf{s}_k^*$ . If  $\phi^{r_i}(\mathbf{s}_k^*) < \phi^{r_i}(\mathbf{s}^*)$  then  $\mathbf{s}^*$  is not an equilibrium of  $\mathcal{G}^*(\mathbf{r})$ , as some agents of type  $i$  would deviate. We conclude that there is at least one type  $i$  s.t.  $r_i < r_k$ , and  $\phi^{r_i}(\mathbf{s}_k^*) \geq \phi^{r_i}(\mathbf{s}^*)$ .

Since  $r_k > r_i \geq 2$ , and  $\phi^{r_i}(\mathbf{s}_k^*) \geq \phi^{r_i}(\mathbf{s}^*)$  by Lemma 6, we get Eq. (1) again. Thus  $SC(\mathbf{s}^*) \leq SC(\mathbf{s}_k^*)$ .

By Lemma 2, we have that  $\text{C-PoA}(\mathcal{G}, \mathbf{r}) \leq \text{C-PoA}(\mathcal{G}, r_k) \leq \alpha_k$  (note that  $\alpha_j = 1$ ).  $\square$

Finally, in a RSG a small fraction of the agents with high uncertainty cannot inflict too much damage.

**Theorem 8.** *Let  $\mathcal{G}$  be an affine RSG. Suppose that  $\mathbf{r}$  is composed of two types,  $r_k > r_j$ . Then  $\text{C-PoA}(\mathcal{G}, \mathbf{r}) \leq \frac{r_j}{2} + \frac{n_k}{n} O(r_k)$ .*

## Worst-Case Regret

For what follows, we assume that  $r_i = r$  for all  $i$ . In addition, we focus on RSGs, as the analysis is non-trivial even for such simple games.

**Equilibrium and convergence.** In a RSG, the set of edges  $E$  is also the set of actions. For every resource  $e$  in state  $\mathbf{s}$ , we have

$$c^{**}(\{e\}, \bar{\mathbf{s}}) = WCR(e, \bar{\mathbf{s}}) = c_e(r\bar{s}_e) - \min_{d \neq e} c_d(\bar{s}_d/r).$$

As  $\mathcal{G}$  is an RSG,  $\mathbf{s} \in EQ(\mathcal{G}^{**}(r))$  if and only if  $WCR(e, \bar{\mathbf{s}})$  is the same for all occupied edges, and at least as high in unoccupied edges. Denote  $MR(\mathbf{s}) = \max_{e \in E: \bar{s}_e > 0} WCR(e, \bar{\mathbf{s}})$ .

Recall that every RSG with player-specific costs has a weak potential (Milchtaich 1996). We show a similar result for an RSG played by WCR players. Our result requires an additional technical property, but allows an explicit construction of the potential. We say that a function  $z(t)$  is  $r$ -convex if  $z'(t/r)/r \leq r \cdot z'(r \cdot t)$  for all  $t$ , where  $c'$  is the derivative of  $c$ . We note that  $r$ -convexity holds for convex and other commonly used functions (see full version).

**Proposition 9.** *Consider an RSG  $\mathcal{G}$ , and suppose that all cost functions are  $r$ -convex. Then  $MR(\mathbf{s})$  is a weak potential function of  $\mathcal{G}^{**}(r)$ . Then if there are WCR moves, there is a WCR move that reduces  $MR(\mathbf{s})$ .*

**Equilibrium quality for Pigou instances.** In the special case of the family of Pigou instances, we have:

**Proposition 10.** *Let  $q_P^{**}(r) = \max_{\mathcal{G}_P} R\text{-PoA}(\mathcal{G}_P, r)$ .*

- (1) *For any  $r \in [1, 2 + \sqrt{3}]$ , we have  $q_P^{**}(r) = \frac{16}{8(r+\frac{1}{r}) - (1+\frac{1}{r})^2}$ .*
- (2) *For any  $r \geq 2 + \sqrt{3}$ , we have  $q_P^{**}(r) = \frac{(r+\frac{1}{r})^2}{8(r+\frac{1}{r}) - 16}$ .*

Thus  $R\text{-PoA}(\mathcal{G}_P, r)$  for  $r \geq 2 + \sqrt{3}$  is an increasing function in  $r$ , and asymptoting to  $r/8 + o(r)$  (see Figure 2). It remains an open question to derive an upper bound for games over more complex networks.

## Discussion

**Related work.** We focus on previous work on uncertainty (especially strict uncertainty) in congestion games (CG). Strict uncertainty been considered by Ashlagi et al. (2006; 2009). They analyze “safety-level” strategies (similar to WCC) for agents who do not know the total number of players  $k$ , but only an upper bound on  $k$ , and focus on proving the existence of a symmetric mixed equilibrium. For atomic RSGs, they show that (as in our case) uncertainty improves the social welfare. Both the the analysis techniques

and the reasoning required from agents in their setting are quite complex, despite the focus on a simple class of games, whereas in our case the game with WCC behavior reduces to a modified (player-specific) congestion game.

The next two papers are closer to our approach, where agents react to some noisy variation of the current state. Meir et al. (2012) study the effect of Bayesian uncertainty due to agents who may fail to play with some probability (in the spirit of trembling-hand perfection). They focus mainly on RSGs and show that the PoA generally improves if failure probabilities are negligible, but not if they are bounded from 0. Angelidakis et al. (2013) study a related model of RSGs with agents who react based on a quantile of the cost distribution. We further discuss these two papers below.

Piliouras et al. (2013) also study CGs with strict uncertainty with motivation similar to ours, and look at a wide range of decision-making approaches including WCC and WCR. However, agents in their model know the actions of others exactly, and uncertainty over costs stems from the unknown *order* in which players arrive. Our model is more direct, and closer to the traditional view of congestion games.

Babaioff et al. (2007) study the effect of introducing a small fraction  $v$  of malicious agents. This is related to our WCC model, where agents behave *as if* the load on each edge is increased by  $v = (r - 1)\bar{s}_e$  additional (malicious) agents. However, in our model this added load only affects the behavior of real agents, and does not affect the outcome directly. Babaioff et al. (2007) observe that adding some amount of malicious agents may decrease the social cost in equilibrium; see also followup work (Blum et al. 2008; Roth 2008).

Lastly, Halpern and Pass (2012) suggested *iterated regret minimization* as an explanation to human behavior in many games. We emphasize that in contrast to our model, such behavior requires agents to explicitly reason about the *intentives* of other agents.

**Distance-based uncertainty.** Our epistemic model adapts the multiplicative distance-based uncertainty model, initially introduced in the context of voting theory (Meir, Lev, and Rosenschein 2014; Meir 2015). There, the state was the number of votes for each candidate, analogous to the measure of agents on each resource in the present setting. While the epistemic model in both papers is derived from a similar approach, both the behavioral heuristics and the techniques for analysis are quite different. For example in voting games payoffs are highly discontinuous, so it is not *a priori* clear whether pure equilibria exist.

There may appear to be a contradiction between our notion of uncertainty and equilibrium play: if agents converge to a particular state and play it repeatedly, then after a while we might expect them to be certain about this state. However even in “standard” congestion game the state (action profile) is only an abstraction of reality, where there is noise from various sources— from players’ actions and failures, to varying costs. Thus an equilibrium is a fixed point in the abstract model, although in reality there remain fluctuations around the equilibrium point. Thus even in equilibrium there may

be some uncertainty about the exact loads.

Some papers model the underlying distribution explicitly (e.g., (Meir et al. 2012; Angelidakis, Fotakis, and Lianas 2013)), and assume a belief structure that is derived from this distribution. Such an approach does not necessarily provide a better description of the way human players perceive the game. In our model we avoid such an explicit description, and instead use  $s$  (as an abstraction of the current state) and derive agents’ beliefs directly from this state using the distance metric. These beliefs may or may not be consistent with the “real” underlying distribution, which may be highly complex. This simple belief structure allows us to derive PoA bounds on a much wider class of games.

Distance-based uncertainty can also be derived from a statistical viewpoint. Suppose that an agent believes that the actual load is distributed around the reference load  $\bar{s}_e$ . A simple heuristic considers a confidence interval around  $\bar{s}_e$ , with the size of the interval modeled through  $r_i$ . Under this interpretation,  $r_i$  is higher for agent types that are either more risk-averse or less-informed.<sup>5</sup> A crucial point is that if agents act independently the actual congestion would be highly concentrated around its expectation (that is,  $r_i$  should approach 1 as the size of the population grows), and for nonatomic agents there should be no uncertainty at all. However, experimental work in behavioral decision making suggests people perceive uncertainty over quantities as if the standard deviation is proportional to the expectation, even when this is false (Kahneman and Tversky 1974; Tversky and Kahneman 1974).<sup>6</sup>

While in our model the uncertainty is over the actions of the other agents, an alternative way is to present it as strict uncertainty over the agent’s own cost function, also known as Knightian Uncertainty (Knight 1921). Chiesa et al. (2014) have recently applied Knightian uncertainty to auctions, where an agent may have a valuation for an item, but only be aware of some *interval* in which this valuation resides. Further studying the conceptual and technical connection between distance-based uncertainty and Knightian uncertainty may help to gain better understanding of both concepts.

## Conclusions

Game-theoretic models should explain and predict the behavior of players in games. Merely adding to such models uncertainty about the environment is insufficient, since as Simon (1957) wrote: “...*the state of information may as well be regarded as a characteristic of the decision-maker as a characteristic of his environment.*” Indeed, psychological studies suggest that people are both risk-averse and avoid

<sup>5</sup>For example, if the noise on edge  $e$  is a normal distribution with standard deviation  $\sigma$  and mean 0, and an agent of type  $i$  requires confidence of 95% (roughly two standard deviations), then this translates to a strict uncertainty interval of  $[\bar{s}_e - 2\sigma, \bar{s}_e + 2\sigma]$

<sup>6</sup>The most famous example is an experiment where subjects are told the average number of girls born daily in a hospital is  $s$ . People believe that the probability that on a given day the number is within  $[(1 - r')s, (1 + r')s]$  is fixed and does not depend on  $s$ . Kahneman and Tversky (1974) highlight the contrast with standard, statistical analysis, where the range  $r'$  is proportional to  $1/\sqrt{s}$ .

probabilistic calculations (Tversky and Kahneman 1974; Slovic, Fischhoff, and Lichtenstein 1980), and raise concerns with standard models of rationality. We believe that our model captures these behavioral assumptions, and that it is simpler than other approaches for uncertainty representation. In addition, we show that the model still permits the use of standard game-theoretic tools such as equilibrium analysis. The model is flexible, and variations of the belief structure (the distance metric) can be easily made. One limitation of our approach is that distance based uncertainty (an interval) cannot capture bi-modal scenarios. For example, when congestion is usually mild, but in rare cases (an accident) congestion is very high. This is similar to the reason that a normal distribution cannot capture such a scenario.

Our results show that in a risk-averse population, an intermediate level of uncertainty helps to align the incentives of the agents with those of the society. This message is emphasized by showing similar results for two different interpretations of risk-aversion, and is consistent with findings from other models of uncertainty (see Related Work).

Finally, lab experiments with routing games show that human subjects converge to states that are close to, but do not coincide with the Nash equilibria (Rapoport et al. 2009) (especially when looking at individual behavior rather than aggregate congestion). We believe that these discrepancies might be at least partially due to bounded-rational behavior of the agents, who may be risk-averse and/or applying simple variations of best-response heuristics. Thus our model might be able to better explain observed outcomes in congestion games. More empirical and experimental work is required in that respect.

## References

- Aashtiani, H. Z., and Magnanti, T. L. 1981. Equilibria on a congested transportation network. *SIAM Journal on Algebraic Discrete Methods* 2(3):213–226.
- Angelidakis, H.; Fotakis, D.; and Lianas, T. 2013. Stochastic congestion games with risk-averse players. In *Algorithmic Game Theory*. Springer. 86–97.
- Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2006. Resource selection games with unknown number of players. In *AAMAS'06*, 819–825.
- Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2009. Two-terminal routing games with unknown active players. *Artificial Intelligence* 173(15):1441–1455.
- Aumann, R. J. 1999. Interactive epistemology I: Knowledge. *International Journal of Game Theory* 28(3):263–300.
- Babaioff, M.; Kleinberg, R.; and Papadimitriou, C. H. 2007. Congestion games with malicious players. In *EC'07*, 103–112.
- Blum, A.; Hajiaghayi, M.; Ligett, K.; and Roth, A. 2008. Regret minimization and the price of total anarchy. In *STOC'08*, 373–382.
- Chiesa, A.; Micali, S.; and Zhu, Z. A. 2014. Knightian self uncertainty in the vcg mechanism for unrestricted combinatorial auctions. In *EC'14*, 619–620.
- Cole, R.; Dodis, Y.; and Roughgarden, T. 2003. How much can taxes help selfish routing? In *EC'03*, 98–107.
- Halpern, J. Y., and Pass, R. 2012. Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74(1):184–207.
- Hyafil, N., and Boutilier, C. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *UAI'04*, 268–277.
- Kahneman, D., and Tversky, A. 1974. Subjective probability: A judgment of representativeness. In *The Concept of Probability in Psychological Experiments*. Springer. 25–48.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin.
- Meir, R.; Tennenholtz, M.; Bachrach, Y.; and Key, P. 2012. Congestion games with agent failures. In *AAAI'12*.
- Meir, R.; Lev, O.; and Rosenschein, J. S. 2014. A local-dominance theory of voting equilibria. In *EC'14*.
- Meir, R. 2015. Plurality voting under uncertainty. In *AAAI'15*. To appear.
- Milchtaich, I. 1996. Congestion games with player-specific payoff functions. *Games and economic behavior* 13(1):111–124.
- Milchtaich, I. 2000. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research* 25(3):349–364.
- Milchtaich, I. 2005. Topological conditions for uniqueness of equilibrium in networks. *Mathematics of Operations Research* 30(1):225–244.
- Piliouras, G.; Nikolova, E.; and Shamma, J. S. 2013. Risk sensitivity of price of anarchy under uncertainty. In *EC'13*, 715–732.
- Rapoport, A.; Kugler, T.; Dugar, S.; and Gisches, E. J. 2009. Choice of routes in congested traffic networks: Experimental tests of the Braess paradox. *Games and Economic Behavior* 65(2):538–571.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1):65–67.
- Roth, A. 2008. The price of malice in linear congestion games. In *Internet and Network Economics*. Springer. 118–125.
- Roughgarden, T., and Tardos, É. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47(2):389–403.
- Roughgarden, T. 2003. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67(2):341–364.
- Roughgarden, T. 2007. Routing games. In et al., N., ed., *Algorithmic game theory*. Cambridge University Press New York.
- Savage, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46(253):55–67.
- Schmeidler, D. 1973. Equilibrium points of nonatomic games. *Journal of Statistical Physics* 7(4):295–300.
- Simon, H. 1957. A behavioral model of rational choice. New York: Wiley.
- Simon, H. A. 1987. Bounded rationality. *The New Palgrave: utility and probability* 15–18.
- Slovic, P.; Fischhoff, B.; and Lichtenstein, S. 1980. Facts and fears: Understanding perceived risk. In *Societal risk assessment*. Springer. 181–216.
- Tversky, A., and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185(4157):1124–1131.
- Wald, A. 1939. Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics* 10(4):299–326.