

Wardrop Equilibria with Risk-Averse Users

Fernando Ordóñez

Department of Industrial and Systems Engineering, University of Southern California,
Los Angeles, California 90089, fordon@usc.edu

Nicolás E. Stier-Moses

Graduate School of Business, Columbia University, New York, New York 10027, stier@gsb.columbia.edu

Network games can be used to model competitive situations in which agents select routes to minimize their cost. Common applications include traffic, telecommunication, and distribution networks. Although traditional network models have assumed that realized costs only depend on congestion, in most applications they also have an uncertain component. We extend the traffic assignment problem first proposed by Wardrop in 1952 by adding random deviations, which are independent of the flow, to the cost functions that model congestion in each arc. We map these uncertainties into a Wardrop equilibrium model with nonadditive path costs. The cost on a path is given by the sum of the congestion on its arcs plus a constant safety margin that represents the agents' risk aversion. First, we prove that an equilibrium for this game always exists and is essentially unique. Then, we introduce three specific equilibrium models that fall within this framework: the *percentile equilibrium* where agents select paths that minimize a specified percentile of the uncertain cost; the *added-variability equilibrium* where agents add a multiple of the variability of the cost of each arc to the expected cost; and the *robust equilibrium* where agents select paths by solving a robust optimization problem that imposes a limit on the number of arcs that can deviate from the mean. The percentile equilibrium is difficult to compute because minimizing a percentile among all paths is computationally hard. Instead, the added-variability and robust Wardrop equilibria can be computed efficiently in practice: The former reduces to a standard Wardrop equilibrium problem and the latter is found using a column generation approach that repeatedly solves robust shortest path problems, which are polynomially solvable. Through computational experiments of some random and some realistic instances, we explore the benefits and trade-offs of the proposed solution concepts. We show that when agents are risk averse, both the robust and added-variability equilibria better approximate percentile equilibria than the classic Wardrop equilibrium.

Key words: wardrop equilibrium under uncertainty; nonadditive cost models; robust game theory

History: Received: April 2007; revisions received: June 2008, March 2009, May 2009; accepted: May 2009.

Published online in *Articles in Advance* January 22, 2010.

1. Introduction

Network games model the interaction between agents who select routes to go from their origins to their destinations. The most common applications can be found in modeling transportation, telecommunication, and logistic systems. Although for some applications decisions may be dictated by a system manager, agents frequently select routes on their own, giving rise to competition for the network resources. It is typically assumed that agents are independent and wish to optimize some individual performance measure—such as utility, delay, cost, or profit—until they all collectively achieve an equilibrium situation in which no agent has any incentive to deviate. In network games—or more generally, in congestion games—the dependence that forces agents to strategize their decisions is manifested through a travel time function (also known as latency, cost, or link performance function, depending on the specific application). In their simplest form, these functions map the flow on a link

to the time needed to traverse it. More complicated versions consider tolls and other arc attributes.

Realistically, however, travel time functions are rarely known exactly prior to making a decision. Agents may have estimates of the travel time based on past experience, but even if all decisions made by agents can be forecasted accurately, there are external factors that make travel times uncertain. A way to resolve this deficiency is to consider an imperfect information game in which agents consider uncertain variations from a nominal travel time function and make a decision taking this uncertainty into account. Regardless of the solution concept used, however, agents will have some degree of a posteriori regret because any outcome is unlikely to be exactly at equilibrium with respect to the realized travel times because agents made their decisions before uncertainty is revealed. We therefore make two postulates about how users behave in network games under uncertainty: (1) Different users estimate the effect

of uncertainty according to their respective attitudes toward risk (2) Solutions whose realized travel times are subject to more variability are less stable, thus further away from an equilibrium.

In this article, we generalize the notion of Wardrop equilibria by proposing a solution concept for network games with uncertain travel times where risk-averse agents pad the expected travel time along paths with a constant safety margin that represents the uncertainty. For this reason, the model belongs to the class of nonadditive equilibrium models (Gabriel and Bernstein 1997). Extensions of results by Aashtiani and Magnanti (1981) imply that this model always possesses an equilibrium which is, in addition, essentially unique when the expected travel time functions are increasing with the use of arcs. Under this model, the risk aversion of agents may lead them to take longer but more reliable routes, instead of short but extremely variable ones, potentially providing solutions that are more stable.

To use this framework, we have to explicitly define the safety margin that is added to each path. Our model assumes that we can decompose travel times in a *nominal* term that depends on decisions made by others, and a flow-independent *error* term that captures random effects not explicitly incorporated in the model. In the context of transportation, deviations in travel time can be caused by accidents, traffic signals, road work, weather, or varying traffic conditions; and in the context of telecommunication networks, they can be caused by malfunctioning equipment, noise, interference, signal degradation, or retransmissions. Although the assumption that the deviation from the nominal value is independent from the flow is reasonable in some situations (e.g., the waiting time for a green light in a traffic signal), it is an approximation of reality in other cases. To put this in perspective, though, commonly used tools such as linear regression make a similar assumption: The error is independent from the magnitude of explanatory variables. Similarly to when one uses a regression model, this is an assumption that needs to be verified empirically for the data set used.

We model the uncertainty in travel times with random variables that represent the deviation from the expected travel time on each arc. However, a solution concept that depends on the distributions of these random variables faces significant practical challenges. The distributions may be difficult to measure, and if estimated with sampling, it is not clear how solutions could be affected by the precision of the sampling procedure. In addition, the estimation procedures may not be realistic from the point of view of agents' capabilities. Another alternative would be to use scenarios to capture uncertainty. However, the combination of scenarios creates a massive number of

cases, leading to computationally intensive methods. Finally, uncertain parameters are sometimes replaced by their expected values to get first-order effects. This approach leads to a solution, henceforth referred to by *nominal equilibrium*, that ignores the uncertainty altogether. It is unlikely that these solutions are good when agents are risk-averse because they consider the variability explicitly when making decisions.

A natural model that captures the risk aversion of agents considers generalized costs that are defined as a specified percentile of travel times along each path. The corresponding solution concept is referred to by *percentile equilibrium*. Even though it is appealing, this model seems to be computationally hard to solve because, even under simple assumptions such as flow-independent i.i.d. random variables in every arc, one may need to enumerate all paths to find the one that minimizes the percentile of the travel time. We therefore propose two other solution concepts based on the same framework and evaluate them in terms of how well they approximate percentile equilibria. These two simpler approaches require estimating a single parameter per arc that measures the maximum additive deviation from the nominal travel time, assumed to be independent of the flow. This is appealing from the point of view of the agent because data requirements are lighter and more in line of what an agent may know in practice: it is not likely that agents know the distribution of travel times exactly although they probably have some idea of the variation present in each link.

The first approximation to percentile equilibria we consider is an equilibrium model proposed by Uchida and Iida (1993) that adds a fixed fraction of the maximum deviation to the nominal value in each arc, henceforth referred to by *added-variability equilibrium*. Computationally, this model has the same complexity as computing a nominal equilibrium. The problem with this solution concept is that it cannot be used when agents have heterogeneous preferences toward risk, or when some origin-destination (OD) pairs are much further away than others because an additive model like this one fails to consider risk-diversification effects.

The second approximation considers that agents react to the uncertainty present in the system by solving a robust optimization problem. Robust optimization has become a popular paradigm in mathematical programming, gaining a wide acceptance in a number of applications such as portfolio optimization, supply chain management, and network design, to name a few. This paradigm addresses optimization problems with uncertain parameters by finding a solution that has optimal worst-case objective. Instead of using distributions, it is considered that the uncertain parameters belong to a bounded convex set. Such

sets can represent the estimation confidence intervals of the uncertain parameters and also model interactions or correlations between them. Intuitively, these sets prohibit all parameters from taking their worst-case values simultaneously because such an event is extremely unlikely. Indeed, although we take a worst-case perspective, to avoid being overly pessimistic, we omit unlikely situations where a large number of links encounter large deviations from their nominal travel times. Another positive feature of this methodology is that for many types of problems and uncertainty sets, solving the robust optimization problem has the same computational complexity as solving the deterministic version of the problem.

In our case, agents solve robust shortest path problems: each agent selects the path that has the best worst-case travel time. Here, the worst-case travel time is computed assuming that the number of arcs that deviate from the nominal value does not exceed the *budget of uncertainty* of the model, and that the deviation is maximal. This budget of uncertainty is a parameter associated with every agent that captures her degree of risk aversion. We call the resulting solution a *robust equilibrium*.

We present a column generation approach for computing robust equilibria that uses the robust shortest path problem to select the paths to be added to the restricted master problem. This approach puts together several pieces existing in the literature: we use the framework of nonlinear complementary problems of Aashtiani and Magnanti (1981), the column generation algorithm for nonadditive travel time functions of Gabriel and Bernstein (1997), and the robust shortest path problem put forward by Bertsimas and Sim (2003). The complexity of the algorithm is exactly the work required to compute a deterministic equilibrium for a problem with nonadditive costs. This means, however, an increase in complexity when the deterministic problem has additive costs. For the important subcase of separable travel time functions, the restricted master problem can be simplified from a nonlinear complementary problem to a convex optimization one.

We perform extensive computational experiments to study how robust and added-variability equilibria compare to each other and how close they are to percentile equilibria. We therefore introduce measures to quantify the distance of a solution to a percentile equilibrium, as well as graphical and statistical tools to help compare the solutions. For instance, we say that a solution is more stable or closer to equilibrium if the distribution of travel time is more concentrated. In addition, we define the regret of an agent as the ratio between the percentile of travel time of her path over the smallest percentile among all paths. Under a percentile equilibrium, all agents experience

a regret equal to one. We provide evidence that under robust and added-variability equilibria, agents have lower regrets than under the nominal counterpart. Hence, when travel times are uncertain and agents are risk averse, the proposed solution concepts may be a better representation of the outcome of the network game.

We structured the paper in the following way. In the remainder of this introductory section we discuss the literature related to this work. Section 2 presents a generic framework for network games that incorporate uncertainty in traversal times of arcs, and shows that equilibria exist and are essentially unique. Section 3 then presents three models that fit the framework introduced earlier. It presents some examples that compare these models and also discusses how to calibrate a model and estimate its parameters. We discuss an efficient column-generation algorithm that finds robust equilibria in §4. In §5, we describe the measures used to compare solutions whereas §6 presents our computational results. Finally, we summarize and indicate further lines of research in §7.

1.1. Related Work

Our model considers network games with an infinite number of infinitesimally small agents and the traditional solution concept of Wardrop equilibria. Wardrop (1952) introduced this game when he postulated that agents in a transportation network select routes of minimal delay with respect to the prevailing conditions. Beckmann, McGuire, and Winsten (1956) were the first to prove that a Wardrop equilibrium always exists and is essentially unique. Because of its simplicity, practitioners have extensively used this model and some of the extensions that have been introduced since its creation. For more details and references in the context of traffic networks, we refer the reader to the book by Sheffi (1985), and in the context of telecommunication networks, to the survey by Altman et al. (2006).

Most of the network models developed to date assume that delays can be predicted accurately. However, it has been recognized that this is not necessarily the case in practice. For example, Liu et al. (2002) develop a dynamic traffic assignment model with stochastic components, building on the work of Mirchandani and Soroush (1987). In their model, decision makers minimize the expected travel time and errors are normally distributed. Furthermore, Uchida and Iida (1993) added a safety margin to expected travel times on each arc. This safety margin is dependent on the degree of risk aversion of agents and proportional to the standard deviation of travel time on the arc. Most related to our framework of robust equilibria, Bell and Cassir (2002) propose a model in which a *demon* tries to maximize agents'

travel time by congesting a worst-possible single arc in the path of the agents. This is related to a robust shortest path when only one arc can deviate. The two approaches are not equivalent, though, because Bell and Cassir consider a game theoretical setting in which the demon plays a mixed strategy randomizing over actions that increase the congestion in one arc at a time; the robust equilibrium that we introduce considers worst-case deviations for every path and an arbitrary number of arcs. Finally, Lo and Tung (2003) introduce a probabilistic user equilibrium model for networks with stochastic capacity. This equilibrium model requires that used routes not only have the same mean travel time value but that its variance is bounded by given performance guarantees. The percentile equilibrium concept differs in that we require that any route with positive flow has a constant percentile of travel time.

The robust optimization approach was introduced to the mathematical programming community by Ben-Tal and Nemirovski (1998) and El-Ghaoui, Oustry, and Lebret (1998). The robust solution for an optimization problem under uncertainty is defined as the solution with the best objective value in its worst-case uncertainty scenario. An attractive feature of a robust solution is that it behaves “well” for all likely uncertainty, in particular the actual realization of the uncertain variables, which is impossible to predict a priori. For many problems, finding the robust solution is not harder than solving the deterministic counterpart (Ben-Tal and Nemirovski 1998). Robust optimization has provided interesting answers to applications on structural design (Ben-Tal and Nemirovski 1997), least-square optimization (El-Ghaoui and Lebret 1997), portfolio optimization problems (Goldfarb and Iyengar 2003; El-Ghaoui, Oks, and Oustry 2003), supply chain management problems (Bertsimas and Thiele 2004; Ben-Tal et al. 2005; Bertsimas and Thiele 2006), and integer programming and network flows (Bertsimas and Sim 2003; Atamtürk and Zhang 2007; Ordóñez and Zhao 2007b).

Game theorists have long considered that agents may not have complete information at the time of making their decisions. The first to explicitly consider incomplete information games was Harsanyi (1967, 1968). In these games, agents are assumed to base their decisions on a probability distribution that models what is unknown to them. Harsanyi’s solution concept, called a Bayesian equilibrium, assumes that agents compute their expected payoffs using these prior distributions. A shortcoming of this model is that it is not obvious how agents can estimate the prior distribution. Holmström and Myerson (1983) refined Bayesian games by considering the case where agents need not know the distribution. Indeed, in an

ex post equilibrium no agent has an incentive to deviate from the selected strategy even after learning the realization of all the uncertainty. Although appealing from the modeling perspective, many games—including the one considered in this paper—generally do not admit equilibria of this type. A few recent papers have explored the application of robust optimization to game theory. Hayashi, Yamashita, and Fukushima (2005) characterize *robust Nash equilibria* in simple games as solutions to a second-order cone complementarity problem. Aghassi and Bertsimas (2006) also consider robust games and prove that robust Nash equilibria always exist. These articles on robust game theory consider finite number of agents and do not concentrate on robust equilibria in network settings. Therefore, their findings on the uniqueness of equilibria and their algorithms to compute them do not directly apply to our games.

In our case, each agent needs to solve a shortest path problem under uncertainty. These problems are typically difficult to solve, especially when errors are correlated or when the objective is more complex than a combination of the expected value and the variance of the travel time. An example of the latter is minimizing the percentile of the travel time, as we are going to consider in §3.1. We refer the reader to Andreatta and Romeo (1988) and Bertsekas and Tsitsiklis (1991) for some classic references, and to Fan, Kalaba, and Moore (2005a, b), Nikolova, Brand, and Karger (2006), and Nie and Wu (2009) for some newer ones. All these algorithms either require a knowledge of the distributions or make strong assumptions on their structure. In the context of robust optimization that we are going to use in §3.3, Bertsimas and Sim (2003) study the *robust shortest path problem*. This problem considers distribution-less instances; the required input for each arc is the maximum deviation from the nominal length. The structure perfectly fits our formulation and, therefore, it will be the subproblem that agents solve in our network game.

Another source of potential uncertainty is on the demand side. Traditionally, this has been modeled with the help of a demand function that links the delay experienced and the demand. The reasoning being that there is a latent demand that will be realized when the utility of traveling is low enough because otherwise agents can use external options like other networks or modes of transportation, or even not traveling at all. There is recent work that models demand uncertainty using game-theoretic techniques. Ashlagi, Monderer, and Tennenholtz (2006) and Ukkusuri and Waller (2010) consider network games in which agents have incomplete information about the demand.

Finally, there are transportation models that consider a different type of uncertainty. *Stochastic user*

equilibrium models, first proposed by Dial (1971) and Daganzo and Sheffi (1977), incorporate uncertain variations in the perception of costs by different agents. Instead of considering that the experienced travel time is unknown to the agents as we do, these models assume that it is known but different agents may extract different utility from it. This is modeled by summing independent and identically distributed perturbations to the nominal travel time of each agent. Moreover, it is assumed that these identical distributions are known. For certain distributions, these equilibria are tractable and can be readily computed; for more details, see Sheffi (1985). This approach and ours complement each other, and one could model the two sources of uncertainty together (see, e.g., Liu et al. 2002).

2. General Wardrop Equilibrium Model with Uncertain Travel Times

This section presents a general, nonadditive equilibrium model that considers uncertain travel times, and proves the existence of this type of equilibrium. The following section proposes three concrete models that specify explicitly how to represent the uncertainty.

We consider a directed graph, $G = (N, A)$, together with a set of OD pairs, $K \subseteq N \times N$. For each terminal pair, $k = (s_k, t_k) \in K$, let \mathcal{P}_k be the set of directed (simple) paths in G from s_k to t_k , and let $d_k > 0$ be the demand rate associated with commodity k . We refer to the set of all paths by $\mathcal{P} := \bigcup_{k \in K} \mathcal{P}_k$. A feasible flow h assigns a nonnegative and possibly fractional value h_P to every path $P \in \mathcal{P}$ such that $\sum_{P \in \mathcal{P}_k} h_P = d_k$ for all $k \in K$. The total flow, or load, along arc $a \in A$ can be easily computed by summing over paths: $f_a := \sum_{Q \in \mathcal{P}: Q \ni a} h_Q$.

The expected travel time along an arc a is given by a (deterministic) load-dependent nominal value. This is measured by a nonseparable function $l_a: \mathbb{R}_{\geq 0}^{|A|} \rightarrow \mathbb{R}_{\geq 0}$, assumed to be positive and continuous. Moreover, we assume that the mapping $(l_a)_{a \in A}$ is strictly monotone.¹ With a slight abuse of notation, we also use $l_a(h) := l_a(f)$ for the arc load f defined above. Travel time functions that depend on the full vector of flows arise frequently in practice, the most common examples in the domain of transportation being two-way streets and multimodal networks, and in the domain of telecommunications being interference. Using the additivity of expectations, the nominal travel time of a path $P \in \mathcal{P}$ under a given flow h is simply $l_P(h) := \sum_{a \in P} l_a(h)$.

We consider a situation in which users take uncertainty into account by adding a *padding* to the

expected travel time of a path. A user that plans on taking path $P \in \mathcal{P}$ when the prevailing flow is h realizes that the travel time may deviate from $l_P(h)$ because of uncertainty. We assume users will add a deterministic value of δ_P to penalize the path for its uncertainty. Hence, the modified travel time of the path P will be $l_P(h) + \delta_P$. The value of δ_P can include considerations about the structure of the path, its uncertainty, and the users' attitude toward risk. For example, if the distribution of travel times on path P were known, δ_P could represent a desired percentile level. The following section provides more details of this interpretation as well as others.

Because expected travel times depend on the congestion level, users compete with each other. Wardrop (1952) postulated that users in a network game select routes of minimal travel time. This concept leads to the so-called Wardrop equilibrium, which is a route pattern in which users do not have an incentive to deviate because all users are assigned to paths that are shortest. Actually, this solution is a Nash equilibrium in the game with an infinite number of users (see de Palma and Nesterov 1998). We generalize this solution concept by incorporating the uncertain travel times to the model. The following definition states that a Wardrop equilibrium with uncertain travel times optimizes the users' objective introduced above for all users simultaneously.

DEFINITION 2.1. A flow h is called a *Wardrop equilibrium with uncertain travel times* if and only if

$$l_P(h) + \delta_P \leq l_Q(h) + \delta_Q$$

for all $P, Q \in \mathcal{P}_k, k \in K$ with $h_P > 0$.

In other words, this is a (regular) Wardrop equilibrium with respect to modified costs $\tilde{l}_P(h) := l_P(h) + \delta_P$. The lack of separability with respect to the sum is typically referred to as a model with "nonadditive costs" (Gabriel and Bernstein 1997).

It is well known that the Wardrop equilibrium condition can be expressed as a nonlinear complementarity problem (NCP) in the space of arc-flows (Aashtiani and Magnanti 1981). For the case of nonadditive costs, it is necessary to consider an NCP that uses a path formulation, as shown in the following proposition. The complementarity condition is the key element of this formulation as it says that paths can only route flow if they are shortest with respect to the users' objective. Under the assumption that travel times are strictly monotone, although one can have multiple equilibria arising from different flow decompositions, the corresponding load on arcs coincide and, thus, also their generalized costs for each OD pair. Hence, equilibria are said to be "essentially unique."

¹ A function $F: D \rightarrow \mathbb{R}^n, D \subset \mathbb{R}^n$ is called strictly monotone on D if, for any two distinct $x \in D$ and $y \in D$, $(x - y)^T (F(x) - F(y)) > 0$.

PROPOSITION 2.2. *A flow h is at equilibrium if and only if it solves*

$$0 \leq h_p \perp l_p(h) + \delta_p - w_k \geq 0$$

for all $P \in \mathcal{P}_k, k \in K, \quad (1)$

where the notation \perp means that at least one of the two constraints on either side must be tight, and the free variable $w_k \in \mathbb{R}$ represents the minimal objective function values for the users' objective. Moreover, an equilibrium always exists and is essentially unique.

The NCP displayed in the previous equation expresses that: (i) flows need to be nonnegative ($0 \leq h_p$), (ii) the inequality $l_p(h) + \delta_p \geq w_k$ guarantees that w_k is a lower bound for the objective of all users corresponding to OD pair $k \in K$, and (iii) there is complementarity condition between not using a path and the path being shortest.²

PROOF OF PROPOSITION 2.2. To prove the equivalence and the existence, we rely on Proposition 4.1 and Theorem 5.4 of Aashtiani and Magnanti (1981). The uniqueness part can be derived along the lines of Theorem 6.2 in the same reference. Indeed, let us assume that two distinct flows h^1 and h^2 are at equilibrium. Then, $h_p^i(l_p(h^i) + \delta_p - w_k^i) = 0$ for all $P \in \mathcal{P}_k, k \in K$, and $i = 1, 2$. Using vector notation for l, δ , and w , the previous equation and the nonnegativeness of all the factors implies that $(h^1 - h^2)^T [(l(h^1) + \delta - \Pi w^1) - (l(h^2) + \delta - \Pi w^2)] \leq 0$, where Π denotes the path-OD pair incidence matrix. Notice that the left-hand side equals $(h^1 - h^2)^T [(l(h^1) - \Pi w^1) - (l(h^2) - \Pi w^2)] = (h^1 - h^2)^T [l(h^1) - l(h^2)] + (h^1 - h^2)^T \Pi[w^2 - w^1]$. Using the arc-path incidence matrix Δ and removing the very last term in the previous equation because both flows h^1 and h^2 satisfy demands, $0 \geq (h^1 - h^2)^T \Delta^T [l(\Delta h^1) - l(\Delta h^2)] = (f^1 - f^2)^T [l(f^1) - l(f^2)]$, where the right-hand side is indexed by arcs. This contradicts the strong monotonicity of l , and implies that $f^1 = f^2$ as we wanted to show.

To conclude, consider arbitrary paths P^i that carry flow under h^i for $i = 1, 2$ for an arbitrary OD pair $k \in K$. Because h^1 is at equilibrium and $h_{P^1} > 0$, $l_{P^1}(h^1) + \delta_{P^1} \leq l_{P^2}(h^1) + \delta_{P^2} = l_{P^2}(h^2) + \delta_{P^2}$. This inequality together with the converse one proves that generalized costs under both equilibria are the same. \square

We note that similar results can be obtained using the more general framework of variational inequalities (Smith 1979; Dafermos 1980) if one uses a path formulation as we have done here. In fact, the NCP

displayed in (1) is equivalent to the following variational inequality:

$$\langle l_p(h) + \delta_p, x_p - h_p \rangle_{P \in \mathcal{P}} \geq 0$$

for any feasible flow $x, \quad (2)$

where $\langle \cdot, \cdot \rangle_{P \in \mathcal{P}}$ denotes the standard inner product in $\mathbb{R}^{|\mathcal{P}|}$.

Finally, in the case of separable functions $l_a(\cdot)$ for which the travel time on an arc $a \in A$ depends only on the flow f_a on the same arc, the equilibrium problem can be formulated as a convex optimization problem. This approach was pioneered by Beckmann, McGuire, and Winsten (1956) and consists of finding an objective function such that its first-order optimality conditions match those that define the equilibrium (see Altman et al. (2006) for a background on potential functions). In this case, this is achieved by minimizing

$$\sum_{a \in A} \int_0^{f_a} l_a(z) dz + \sum_{P \in \mathcal{P}} \delta_P h_P \quad (3)$$

over the space of feasible flows. Notice that (3) includes both path- and arc-variables, which makes the mathematical program exponentially large in the worst case. The optimal solution of this convex program verifies Definition 2.1, which implies that an equilibrium always exists. Moreover, if travel time functions are strictly increasing, the problem is strictly convex and hence there is a unique equilibrium in the space of arc-flows (although there may be multiple decompositions into path-flows).

3. Specific Wardrop Equilibrium Models with Uncertain Travel Times

In this section we present three equilibrium models that allow for uncertain travel times and risk-averse users. To that effect, we incorporate explicit representations of uncertainty into the model introduced in the previous section. The following three equilibrium models provide solutions that incorporate uncertainty in different ways. When users are risk-neutral, they prefer to minimize expected costs. In that case, our three solution concepts coincide with a *nominal equilibrium*, which ignores the travel time variability. The three models also coincide in the case of extremely risk-averse users, because in that situation users consider the maximum possible travel times for all arcs and ignore the uncertainty as well. Between those extremes, the solutions may be different, as illustrated by the examples in §3.4.

Although users have arbitrary risk tolerances, we consider that all users that belong to one OD pair are homogeneous. This is without loss of generality,

² Concretely, $h_p \perp l_p(h) + \delta_p - w_k$ means that when $h_p > 0$ then $l_p(h) + \delta_p - w_k = 0$, and when $l_p(h) + \delta_p - w_k > 0$ then $h_p = 0$. Hence, used paths have generalized cost equal to w_k and paths with more cost than w_k are not used.

because we can group users in copies of the same OD pair, according to their risk preferences. We will consider users that are risk-neutral or risk-averse; risk-seeking behavior could be also handled following similar methods, but we do not do it here for simplicity.

It is important to highlight that all users will experience the same travel time, as opposed to the assumptions of stochastic user equilibrium models, in which each user extracts a different utility from the same travel time. In practice, this is implemented with a different realization of a random variable for each user.

Our model of uncertainty represents the deviation from the nominal travel time in each arc a as $u_a Z_a$. Here, Z_a is a random variable with expected value equal to zero and support in $[-1, 1]$, and $u_a \geq 0$ is an upper bound on the maximum possible deviation from the nominal value.³ Although we make no assumptions on the distribution or its independence from other distributions, the deviation on each arc is assumed to be independent of the flow on the network. This assumption fits some situations well and is an approximation of reality in other cases (see §3.5 for a discussion on modeling assumptions). Overall, the actual travel time experienced by all users that selected arc a equals $l_a(h) + z_a u_a$, where z_a is a realization drawn from the distribution Z_a .

3.1. Percentile Equilibrium

One way to capture the risk tolerance of users is to assume that they want to ensure that they arrive at their destination on time. A risk-neutral user minimizes the expected travel time, but a risk-averse user prefers to consider a margin of error. Hence, she is more likely to minimize a higher-than-50th percentile of the experienced travel time. To capture this behavioral assumption with the deterministic model introduced in the previous section, we will let the padding for paths be a percentile of the travel time. Indeed, users of OD pair k traveling along a path $P \in \mathcal{P}_k$ will set δ_P to a percentile of the total deviation from the nominal travel time along P . The percentile level is chosen according to the risk tolerance of the user; risk-neutral users will choose the 50th whereas risk-averse ones may, for example, choose the 90th. Behaviorally, although users are unlikely to know the distributions of travel times in each arc, they may have enough experience in the network to determine the likelihood of making it on time for each possible route. This leads us to the following definition, which is based on Definition 2.1.

³ The discussion below also assumes that Z_a is symmetric so the 50th percentile, that is, the median equals the mean, which is zero. This property is only used to claim that a percentile equilibrium equals a nominal equilibrium, but it is otherwise not required.

DEFINITION 3.1. An α -percentile equilibrium is a Wardrop equilibrium with uncertain travel times where the padding along path $P \in \mathcal{P}_k$ is set to $\delta_P = G_P^{-1}(\alpha)$ and G_P is the cumulative distribution function of the deviation along path P given by the random variable $\sum_{a \in P} u_a Z_a$.

Under a percentile equilibrium, all users follow paths that minimize the α -th percentile of travel time. The results in §2 show that a percentile equilibrium exists and is essentially unique. We note that although users will experience a posteriori regret for any realization of the uncertainty, if the description of the uncertainty is accurate, the empirical α -th percentile of the travel time on every used path $P \in \mathcal{P}_k$ will approach w_k after enough time has elapsed. In other words, even though the paths that are used may have different distributions of travel time, users will experience the same percentile on the long run, causing no long-term regret.

Computing this equilibrium, however, is considerably difficult and not only because (1), (2), and (3) consider path variables. To compute the percentile one needs the full distributions along paths; this calculation cannot be separated by arcs. If one assumes that deviations in different arcs are independent of each other, the distribution of the deviation on a given path is equal to the convolution of all the random variables along that path. Furthermore, computing a shortest path with respect to percentiles is not easy even under simplifying assumptions (such as with normally distributed errors), as we explain in §1.1. This difficulty makes developing an efficient column generation procedure to compute the percentile equilibrium a challenging task.

The next two sections propose computationally-tractable models that are aimed at approximating percentile equilibria, because at least in principle it should provide no regret for risk-averse users. As we shall describe in §5, one of our criteria to evaluate flows will be their similarity to a percentile equilibrium.

3.2. Wardrop Equilibrium with Added Variability

Uchida and Iida (1993) considered adding a safety margin to the expected travel times on arcs to account for user risk aversion. Although the standard deviation of the travel time along a path does not separate by arcs,⁴ this approach can be used as a crude representation of the uncertainty of travel times in the network. Percentiles of travel times along paths do not separate by arcs and are not additive either, but we

⁴ The standard deviation of the travel time along path is not additive; in the case of independent errors, it equals the square root of the sum of the variances of the travel times of the arcs along that path. For this reason, the standard deviation of the sum of variables is less than the sum of the standard deviations, situation normally referred to as the *risk diversification effect*.

shall see that solutions computed with this approach provide reasonable approximations of percentile equilibria in some practical situations. Besides not taking into account the risk diversification effect, the main difficulty that prevents this method from being precise is that it is hard to incorporate arbitrary risk tolerances for different users.⁵ From the point of view of information availability, a big benefit is that this method does not require users to know the whole distribution of the travel times along paths; it is enough that users know an estimate of the variability along each arc. The formal definition of equilibria is based on Definition 2.1 again.

DEFINITION 3.2. A ϕ -added-variability equilibrium is a Wardrop equilibrium with uncertain travel times where the padding along path $P \in \mathcal{P}_k$ is set to $\delta_P = \sum_{a \in P} \phi u_a$.

Here, the choice of the constant ϕ depends on the degree of risk aversion of users. When users are risk neutral, one should choose zero, whereas for extremely risk-averse users, ϕ will be close to one. From a computational perspective, computing these equilibria has the same complexity as computing an equilibrium that ignores uncertainty, because it is a standard (deterministic) Wardrop equilibrium with respect to nominal travel times plus a constant fraction of u_a .

3.3. Robust Wardrop Equilibrium

Now, we introduce another model to approximate the behavior of users that make decisions under uncertain travel times. This model is more tractable than computing a percentile equilibrium but more complex than adding a fraction of the variability to all arcs. The novel element about this approach is that users use robust optimization to select *robust shortest paths*. This means that users take a worst-case perspective without being overly pessimistic and seek a shortest path considering a reasonable estimate for the maximum deviation of the travel time along paths. Because it is unlikely that users face an extreme deviation from nominal travel times in too many arcs, we give users of OD pair k an uncertainty budget of Γ_k . We impose that the relative deviation from the nominal travel time is less than Γ_k on any path. The budget of uncertainty is meant to estimate the number of arcs in which travel times can significantly deviate from nominal values. Important factors that influence the selection of Γ_k are the risk preferences of users of that OD pair and the average path length for that OD pair.

⁵ A possible way to implement arbitrary risk tolerances is to create a copy of the network for each OD pair and to redefine costs to depend on the total flow along all of the copies of the arcs. We do not consider this possibility in this article because the computational study only considers separable instances.

For instance, the case of $\Gamma_k = 0$ corresponds to risk-neutral users, because they ignore uncertainty and only consider the expected travel time. On the other extreme, for extremely risk-averse users, we can consider large values of Γ_k that will make the model use the worst-case realization of uncertainty on every arc.

Mathematically, a user traveling between OD pair k decides the route by selecting the path that has the best worst-case travel time. This amounts to solving the following optimization problem:

$$\begin{aligned} \min_{P \in \mathcal{P}_k} \max_z \left\{ \sum_{a \in P} (l_a(h) + z_a u_a) : \sum_{a \in P} |z_a| \leq \Gamma_k, -1 \leq z_a \leq 1 \right\} \\ = \min_{P \in \mathcal{P}_k} \max_z \left\{ \sum_{a \in P} (l_a(h) + z_a u_a) : \right. \\ \left. \sum_{a \in P} z_a \leq \Gamma_k, 0 \leq z_a \leq 1 \right\}, \quad (4) \end{aligned}$$

where h is the flow that encodes the collective decisions made by all users. Notice that the robust perspective assumes that each element z_a is a decision variable and not a realization of random variable Z_a as previously. Definition 3.3 says that a Wardrop equilibrium is robust when all flow goes along robust shortest paths.

DEFINITION 3.3. A Γ -robust Wardrop equilibrium (RWE) is a Wardrop equilibrium with uncertain travel times where the padding along path $P \in \mathcal{P}_k$ is set to

$$\delta_P := \max \left\{ \sum_{a \in P} z_a u_a : \sum_{a \in P} z_a \leq \Gamma_k, 0 \leq z_a \leq 1 \right\}.$$

As with the previous two equilibria, all results in §2 apply. Hence, a robust equilibrium exists and is essentially unique. Note that deviations from nominal travel times are measured in relative terms because absolute deviations would lead to $\delta_P = \min\{\Gamma_k, \sum_{a \in P} u_a\}$. In that case, we either add the uncertainty budget or add all uncertainties to the nominal travel times along all paths. Both alternatives seem to be of limited interest because users do not end up using the information provided by the uncertainty to select paths.

Robust equilibria and the solution method that we describe can also accommodate more complicated forms of travel-time uncertainty as long as it depends only on the path and is not influenced by the flow. For example, if the uncertain vectors belong to an ellipsoid $\{z: \sum_{a, b \in A} z_a z_b Q_{ab} \leq \Gamma_k^2\}$, then the worst-case deviation on a path, as given by Definition 3.3, has a closed form solution with optimal value equal to $\delta_P = \Gamma_k \sqrt{\phi_P^T (Q|_P)^{-1} \phi_P}$, where $\phi_P = (u_a)_{a \in P}$ and $Q|_P$ is the submatrix corresponding to path P .

After presenting a few examples that fit our model of uncertainty, §4 discusses how such an equilibrium

can be computed efficiently, and later on we study how these three equilibrium notions compare with each other.

3.4. Some Illustrative Examples

In this section we provide two examples that illustrate the three equilibria under uncertainty defined earlier. We start with a concrete application that fits the presented framework and then provide a stylized example that will allow us to discuss some of the details of the model.

While driving in cities, an important source of uncertainty is waiting at traffic lights. The time until a light changes to green can be modeled as a uniformly distributed random variable. Notice that in this example the waiting time is independent from the flow because the random component is how long to wait in the first cycle; the wait time associated with multiple full cycles is not as uncertain and thus is incorporated in the nominal travel time function. Under this example, a percentile equilibrium considers percentiles of the total wait time during red lights, an added-variability equilibrium considers a deterministic time for each traffic light equal to the expected wait plus a fraction of its total time in red, and a robust equilibrium considers that Γ_k traffic lights will have changed to red a moment before the arrival whereas the rest are green. Here, the parameter Γ_k can be calibrated using that the number of red lights can be approximated with a binomial distribution, if we ignore green waves and correlations between cycles times. This binomial would have the number of trials equal the number of traffic lights on paths of OD pair k and the probability equal to time in red over the cycle time.

The following example consists of a simple instance that is related to the classic Braess' paradox network (Braess 2005). The network has a single OD pair (s, t) , connected by a direct arc, by a path with two

intermediate nodes (s, a, b, t) , and by another two paths (s, a, t) and (s, b, t) , each of which skips one node (see Figure 1 where the four nodes are located on the horizontal line). The direct path (s, t) corresponds to a safe (i.e., zero variability) but long highway between s and t . The other paths are composed of arcs that represent local streets. Deviations from nominal travel times in them are because of events such as cabs stopping to pick up or drop off passengers, cars that double park, deliveries to stores, emergency vehicles, etc. Suppose that the uncertainty along paths composed of local streets follows a uniform random variable with support in $[-\beta, \beta]$. This assumption fits the model introduced in §2 and makes it easy to determine the percentile equilibrium. Because users are risk averse, we are interested in a 90th-percentile equilibrium. To compute it, we simply set $\delta_p = 0.8\beta$ for paths composed of local streets, and $\delta_p = 0$ for the path consisting of only arc (s, t) .

Figure 1 shows the three types of equilibrium for this instance. A nominal equilibrium does not make use of the highway because it is too long, but a 90th-percentile equilibrium makes use of all four paths because the risk reduction justifies the extra distance to be traveled. To compute robust and added-variability equilibrium, we need to create a model of variability for arcs. A possibility is to set u_a to a positive value for arcs with uncertainty (local streets) and set $u_{(s,t)} = 0$ (the highway). A reasonable choice is to use the percentile of the travel time for the path on the local streets that compose it. Indeed, in this case the robust equilibrium for $\Gamma = 1$ coincides with the percentile equilibrium because the resulting δ_p is that of the percentile equilibrium in all paths. Added-variability equilibria progressively shift flow from the local streets to the highway when ϕ increases. Doing




Figure 1 Equilibria for the Braess-Like Network

Note. Line widths are proportional to flows. Top left: nominal equilibrium (and 50th-percentile equilibrium). Top center: 90th-percentile equilibrium. Top right: robust equilibria for $\Gamma = 1, 2, 3$. Bottom: added-variability equilibria for $\phi = 0.25, 0.32, 0.375, 0.5, 0.625$.

binary search, we find that the value of ϕ that makes an added-variability equilibrium most similar to the 90th-percentile equilibrium is 0.32. Even in that case, both solutions do not coincide.

Admittedly this example was constructed to illustrate that in some cases a robust equilibrium can achieve what an added-variability equilibrium cannot. Had we modeled uncertainties by setting u_a as specified only in arcs incident to s and to zero otherwise, the added-variability equilibrium would have also coincided with the percentile equilibrium. As with many other modeling problems, there are decisions to be made and the accuracy of the model crucially depends on these decisions. The following subsection further discusses estimation and modeling issues that can arise with this framework in applications.

3.5. Remarks on Calibration of Parameters and Network Representation

In this section, we discuss some of the modeling decisions that have to be made to use the proposed framework. Given a network representation, the travel time function for each arc is estimated according to the geometry of the road (Bureau of Public Roads 1964). Travel time deviations from the nominal values are reasonably easy to estimate by traffic authorities because data sets that contain this kind of information are routinely collected. This data could allow a modeler to calibrate random variables Z_a for each arc in the network. Estimating the travel time deviations over paths, a necessity for computing a percentile equilibrium, is considerably more difficult. A modeler would have to consider the convolutions of arc deviations over paths, possibly adjusted with the correlations between different arcs, or even distributions for end-to-end delay. Obviously, the latter is much harder to collect but many emerging technologies such as global positioning systems and mobile phone data are promising alternatives.

A standard network representation would use arcs and nodes to encode paths that serve demands. However, the same network could be represented by alternative encodings by simply subdividing an arc into multiple parts. Percentile and nominal equilibria are not affected by such a subdivision because the statistics of the complete arc are unaffected. For this reason, a modeler may have no reason to include a node with exactly two incident arcs. However, in some cases added-variability and robust equilibria are influenced by these artificial subdivisions of arcs and by how the variability parameter u_a is set in the different components that form the original arc. Indeed, although maximum deviations are additive so—in theory—subdivisions do not affect added-variability equilibria, a modeler would approximate the maximum deviation by a well-chosen percentile of the distribu-

tion of travel times in that arc. Because an added-variability equilibrium ignores the risk-diversification effect along paths, having more arcs makes this problem more severe. For the case of robust equilibria, having more arcs will make the uncertainty less extreme because the worst-case deviations will be diluted among the additional subdivisions. In summary, both of these models can be adversely influenced by adding artificial subdivisions. Our recommendation is to start with a representation of the network that does not include subdivisions. After evaluating the quality of the equilibrium solution, a modeler could try to improve it by using subdivisions in certain arcs. For example, highways are typically represented by arcs that are much longer than local roads. This unbalance can lead to underestimating the load along highways because under robust equilibria, risk-averse users are not likely to take a highway with a large worst-case deviation. Adding subdivisions provides a mechanism to reverse that effect and improve the fit with the measured flows.

Estimating the parameter Γ_k for the robust equilibrium model presents its own challenges, which require the modeler to strike a balance between risk aversion, the average path length, and how arcs are subdivided. This parameter is harder to estimate than the values that depend only on the network topology because it also includes behavioral aspects. For example, Γ_k could be defined as a fraction of the average number of arcs in paths connecting OD pair k , where the fraction would depend on the characteristics of the uncertainty and risk attitudes. Such a definition should be complemented with a calibration of how users react to uncertainty (see, e.g., Noland et al. 1998; de Palma and Picard 2005). Using simulation a modeler can test if the selected Γ_k gives deviation estimates that match the risk attitudes of the agents.

Let us link this discussion back to the examples of §3.4. In the example concerning traffic lights, travel time variabilities represented times until a green light. In that case, it makes sense to include a node where traffic lights exist, and to map the corresponding distributions of waiting times to those nodes. Hence, there is no reason to add artificial subdivisions. To select values of Γ_k , recall the observation that the number of red lights encountered follows a binomial distribution. In the second example, we saw that the choice of u_a in each arc had an impact of the accuracy of the added-variability equilibrium.

To conclude, it is important to highlight that models usually require refinements to adjust and calibrate their parameters, much as it is done when estimating OD matrices (Sheffi 1985). Our framework makes it relatively easy to compute solutions and their corresponding measures that quantify their resemblance to a percentile equilibrium (see §5). If the approximation

is poor, the modeler can iterate and adjust the different elements until outcomes are acceptable. Section 6 provides an illustration of choosing solutions where most users have low regret, meaning that most of them are satisfied with the route they have chosen.

4. Column Generation Algorithm for Computing RWE

This section focuses in the computation of robust Wardrop equilibria (RWE). A central difficulty is that the users' objective function is not separable. Hence, arc formulations are unsuitable and we have to resort to path variables to solve the NCP shown in (1). This section outlines an algorithm that finds a RWE with a column generation scheme. To generate new columns (variables), the algorithm solves a robust shortest path problem at each iteration.

We follow the approach of Gabriel and Bernstein (1997), who studied a column generation algorithm for solving NCPs arising from network equilibrium problems with nonadditive costs. In each iteration, the column generation algorithm considers only a subset of the possible paths and solves a restricted version of the equilibrium problem. New paths are incorporated as needed. For background and references on column generation and simplicial decomposition algorithms with a focus on network equilibria, we refer the reader to Patriksson (1994) and the references therein.

The algorithm maintains an *active* set of paths $\mathcal{P}' = \bigcup_{k \in K} \mathcal{P}'_k$ that it works with, where $\mathcal{P}'_k \subseteq \mathcal{P}_k$. At each step the algorithm solves the following NCP, also referred to as the *restricted master problem*, to get an equilibrium solution that only takes into account the paths in \mathcal{P}' . In other words, we find a flow h and minimum travel time vector w that solve

$$0 \leq h_p \perp l_p(h) + \delta_p - w_k \geq 0 \quad \text{for all } P \in \mathcal{P}'_k, k \in K. \quad (5)$$

Methods for solving this NCP are readily available in the literature; see, e.g., Patriksson (1994), Nagurney (1999), and Ferris, Mangasarian, and Pang (2001).

The termination condition for the column generation algorithm is to check that the equilibrium solution over the set \mathcal{P}' given by the flow h is indeed at equilibrium over the complete set of paths \mathcal{P} . This is verified when there is no path in $\mathcal{P} \setminus \mathcal{P}'$ that is shorter than the corresponding generalized travel time w_k for the objective induced by h . In other words, given a flow h , we solve the following shortest path problem with nonadditive costs (but linear when formulated in terms of paths):

$$v(h) := \min \left\{ \sum_{P \in \mathcal{P}} (l_P(h) + \delta_P) x_P : \sum_{P \in \mathcal{P}_k} x_P = d_k \text{ for } k \in K, x \geq 0 \right\}. \quad (6)$$

PROPOSITION 4.1. *A flow h^* is a RWE if and only if $v(h^*) = \sum_{P \in \mathcal{P}} (l_P(h^*) + \delta_P) h^*_P$.*

PROOF. The flow h^* solves (6) for $h = h^*$ if and only if h^* is the solution to the VI in (2). \square

Algorithm 1 outlines the major steps of the column generation procedure. This algorithm finishes in a finite number of iterations because in the process of checking whether the candidate flow h^* is a global solution, we identify at least one new path to add to the active set \mathcal{P}' or we finish. Hence, in the worst case we iterate until $\mathcal{P}' = \mathcal{P}$, which makes the problems in Equations (1) and (5) identical. Note that in practice the column generation algorithm converges in a small number of iterations, and it need not enumerate all paths, as it succeeds in quickly identifying the typically small number of paths used in the equilibrium solution.

Algorithm 1 (Column generation)

- 1: Initialize: Add arbitrary paths to $\mathcal{P}'_1, \dots, \mathcal{P}'_K$.
- 2: Set $h^* = 0$ and $v(h^*) = -\infty$.
- 3: **while** $v(h^*) < \sum_{P \in \mathcal{P}'} (l_P(h^*) + \delta_P) h^*_P$ **do**
- 4: Solve the restricted master NCP in (5).
 Let h^* be the optimal solution.
- 5: Solve Problem (6). Let x^* be the optimal solution and $v(h^*)$ be its value.
- 6: Add paths used in x^* to \mathcal{P}' .
- 7: **stop.** The flow h^* is a RWE.

4.1. Robust Shortest Path Problem

Gabriel and Bernstein (1997) point out that solving a shortest path problem with nonadditive costs is a central step in the column generation algorithm for network equilibria. However, they simply mention that the shortest path problem can be solved efficiently because of the structure of the travel time function without describing the procedure in detail. In this section, we outline a procedure to solve the shortest path problem for computing a RWE.

For a fixed flow h , (6) has a linear objective function. As there are no capacity constraints, we can solve this problem by separating it by commodity. Therefore, we can assume that there is an optimal solution that sets $x_P = d_k$ for exactly one path $P \in \mathcal{P}_k$ for each $k \in K$. What is left is to find the path $P \in \mathcal{P}_k$ with smallest $l_P(h) + \delta_P$ for each $k \in K$. The major difficulty of this problem is the exponential number of paths. We avoid this complication by going back to arc variables and explicitly representing δ_P by a maximization of additional auxiliary variables per arc. For each $k \in K$, let us consider the feasible region $X_k = \{x \in \{0, 1\}^A : \sum_{a \in N^+(i)} x_a - \sum_{a \in N^-(i)} x_a = r_i \text{ for } i \in N\}$, where

$r_{sk} = 1$, $r_{tk} = -1$, and $r_i = 0$ otherwise. The problem in question is

$$v_k(h) := \min_{x \in X_k} \left\{ \max_z \left\{ \sum_{a \in A} (l_a(h) + z_a u_a) x_a : \sum_{a \in A} z_a \leq \Gamma_k, \right. \right. \\ \left. \left. 0 \leq z_a \leq 1 \text{ for all } a \in A \right\} \right\}. \quad (7)$$

As the maximum computes $l_p(h) + \delta_p$ for the path identified by x , we have that $v(h) = \sum_{k \in K} d_k v_k(h)$. Note that the budget of uncertainty Γ_k in this problem could be distributed across all arcs, whereas in the robust network equilibrium model and in (6), only deviations in the selected path are considered. Nevertheless, (6) and (7) are equivalent because of the integrality of x .

The problem represented in (7) is called the *robust shortest path problem*, which was introduced by Bertsimas and Sim (2003). As they suggested, by taking the linear programming dual, one can formulate it as the following mixed integer program:

$$v_k(h) = \min_{x \in X_k, y \geq 0, \theta \geq 0} \left\{ \sum_{a \in A} l_a(h) x_a + y_a + \theta \Gamma_k : \right. \\ \left. \theta + y_a \geq u_a x_a \text{ for } a \in A \right\}. \quad (8)$$

In addition, Bertsimas and Sim provide an algorithm that solves the problem with $|A|$ calls to a regular shortest path routine with modified linear cost functions. Indeed, the optimal solution for (8) satisfies $y_a = \max\{u_a x_a - \theta, 0\} = \max\{u_a - \theta, 0\} x_a$ because $x_a \in \{0, 1\}$. Substituting this in (8), it can be shown that the optimal θ will equal u_a for some $a \in A$. With that, an optimal shortest path is simply the minimum among all shortest paths after solving for all $\theta = u_a$ with $a \in A$.

5. Evaluation of Wardrop Equilibria with Uncertain Travel Times

Because travel times are uncertain and unknown to users at the time when they make their routing decisions, none of the equilibrium concepts that we consider in this paper are necessarily at equilibrium with respect to the realized travel times. Indeed, realized travel times for different users corresponding to an OD pair are typically not constant, as would be the case in an actual equilibrium. We will, thus, consider that an outcome is better when more users take paths that they do not regret having taken. In addition, we will also look at the variability of realized travel times, which we will refer to as the unfairness. We shall evaluate the solutions provided by the different equilibrium concepts, and compare the summary statistics of regret, unfairness, and travel times.

Inspired by the work of Roughgarden (2002) and Jahn et al. (2005), we define the α -percentile regret, or simply regret, of a user traveling along path $P \in \mathcal{P}_k$ to be the ratio of her α -percentile of travel time along P to the minimum α -percentile over all $Q \in \mathcal{P}_k$.⁶ This definition allows us to consider the system as a whole because different OD pairs become comparable. At the system level, we get a distribution of regret coming from the values experienced by different users. Under an exact α -percentile equilibrium, the distribution of α -percentile regret collapses to the value one with probability one because all users travel along paths with minimal α -percentile of travel time. For example, assuming that errors Z_a have symmetric distributions, the 50th-percentile regret equals one for all users under a nominal equilibrium. When comparing two solutions, we say that one is better if more users experience less regret than in the other. Later on, in our computational study, we will evaluate these distributions for the different proposed solutions and different values of α .

Besides looking at regret, we also aggregate the distribution of travel times along a path to all the users of one OD pair. In this way, we can complement the previous measure with the study of the intrinsic variability of a solution. Under an equilibrium with respect to experienced travel times, all users would experience equal travel times. In practice, this equilibrium is not implementable because users do not have the ability to forecast travel times precisely. Following the postulates for user behavior, a solution with smaller variability will be considered closer to being at equilibrium and provides another desirable objective for a solution. This variability can be quantified using the standard deviation, the interquartile range, or can be estimated directly from the distributions of travel time. Because this type of analysis is done for each OD pair, we will only do it for instances with a small number of OD pairs.

It is important to highlight that comparing magnitudes of travel times or their expected values across different equilibria cannot be used to provide support for one equilibrium or another because users are competing against each other and they are not centrally controlled.⁷ We believe that users will gravitate toward the more stable solution, which is the

⁶ Another possibility for the denominator is to use the minimum α -percentile of travel time among *used* paths of OD pair k , instead of among all paths. In this case, the ratio measures unfairness among different users. Our choice is based on the available options to a user but, in any case, outcomes are not significantly affected because all of these values are highly correlated.

⁷ For example, consider the classic Braess' paradox network (Braess 2005). At equilibrium, users take long paths although the system (and the users) would be better off if all users took alternative paths. The problem is that the alternative is not stable because users would regret having taken the "good" path and would gravitate

solution with less variability in travel times. In any case, we also provide travel times statistics for each solution that quantify their efficiency. The more inefficient solutions are, the more incentive there is for the system manager to try to encourage cooperative behavior (e.g., by using some form of congestion pricing).

6. Computational Results

The main goal of this section is to show how the three models described in §3 relate to each other. To that extent, we present computational experiments that evaluate nominal, added-variability, and robust equilibria and compare how close they are to being percentile equilibria. We do not compute percentile equilibria explicitly because that would require evaluating all routes in the network and that is prohibitive for networks with more than a few nodes. Although not a requirement of the algorithm that computes robust equilibria, we use a single uncertainty budget Γ for all OD pairs to simplify the presentation and to make it easier to provide insights from the computational results.

Table 1 describes the instances solved by this computational study. We note that we can solve these instances, which range from small to moderately sized, without difficulties. We first conduct experiments on small artificial instances to generate intuition. Next, we verify these findings on the Sioux Fall and Friedrichshain instances, which are larger instances inspired from real-world networks (Bar-Gera 2002, Jahn et al. 2005).⁸ Note that all these instances make use of separable travel time functions, which allows us to compute equilibrium flows relying on the convex optimization formulation of (3), instead of having to solve NCPs. The travel time functions in all these instances are defined by the widely-used Bureau of Public Roads (BPR) functions (Bureau of Public Roads 1964) which equal

$$l_a(f_a) := t_a \left(1 + \beta_a \left(\frac{f_a}{c_a} \right)^4 \right). \quad (9)$$

Here, t_a is the free-flow travel time, c_a is the capacity and β_a is the congestion factor corresponding to arc a .

We perform the computational study using the algorithms described in §4 (the nominal and added-variability equilibria are special cases). We use AMPL

back to the inefficient solution. In other words, the only solution in which the regret for all users equals one is the stable but inefficient equilibrium.

⁸ The instances used are available in the authors' Web pages (<http://www.columbia.edu/~ns2224>). In the interest of space we present here a subset of the computational results obtained. We created an online supplement to this article, also located in the authors' webpages, that contains additional solutions and figures.

Table 1 Problem Instances Used in the Computational Study

Instance name	Short name		V	A	K	A · K	Trials
	Source						
Grid network A, B, and C	Grid.	Ours	24	38	1	38	2,000
Sioux falls (simplified)	SfS	Ours	24	76	5	380	2,000
Sioux falls (complete)	SfC	TNTP	24	76	528	40 K	100
Friedrichshain	Fri	JMSS	224	523	506	265 K	300

Note. TNTP refers to an instance repository maintained by Bar-Gera (2002); JMSS refers to a paper by Jahn et al. (2005).

(Fourer, Gay, and Kernighan 2002) to implement the column generation procedure described by Algorithm 1 and LOQO (Vanderbei 1999) to solve the restricted version of (3). For the robust shortest path problem, we solve the mixed integer program (8) with CPLEX (ILOG 2005). The bottleneck of the computation is performing the following simulation procedure and to a lesser extent solving the restricted master problem so we do not expect significant changes in the running times had we implemented a more efficient algorithm to solve the robust shortest path problem.

To evaluate the descriptive statistics of regret, unfairness, and travel times, we perform a Monte Carlo simulation of the network. To get the realizations of deviations, we assume that the random variables Z_a are uniformly distributed on its support $[-1, 1]$, independent of the uncertainty on other arcs. The assumption of a uniform distribution is made for simplicity; we have adopted some other distributions for some limited tests without a qualitative change in the results.

For a given solution h , we obtain empirical distributions of travel times experienced by users by repeatedly drawing random numbers z_a from Z_a that represent the actual travel time deviations in that trial. A user that selected a path $P \in \mathcal{P}_k$ experiences a travel time of $\sum_{a \in P} (l_a(h) + z_a u_a)$. These realizations of travel times allow us to compute good estimates of percentiles of travel times for each path. Computing the ratio between percentiles and the smallest among them, we get the the distributions of percentile regret.⁹ In addition, we also get the distribution of travel times for each OD pair by considering also that a user selects a given path $P \in \mathcal{P}_k$ with probability h_p/d_k . The column "trials" in Table 1 shows the number of repetitions done for each instance. We performed a sensitivity analysis to determine the number of trials that delivered accurate estimates without making the simulations overly slow. For the

⁹ In practice, we can only estimate regrets approximately because estimating them exactly requires enumerating all paths, which is not tractable. We develop a heuristic that enumerates the paths that are likely to minimize the α -percentile of travel time.




Figure 2 Normal and Robust Equilibria of Instance Grid A

Notes. Left: Worst-case deviations u_a . Center: Nominal equilibrium. Right: Robust equilibria for $\Gamma = 1, 3, 5$.

small instances we can quickly do 2,000 repetitions. For instances SfC and Fri, we obtained an almost constant 50th-percentile regret under a nominal equilibrium after the number of repetitions mentioned in the table.

6.1. Grid Instances

Our first set of experiments studies six by four grid networks with randomly generated uncertainty parameters. We use small instances to describe the methodology in more detail and to derive some insights for the different equilibria that will carry over to the more realistic instances presented later.

To study the effect of travel time variabilities, we set all arcs equal in expectation but with random uncertainty parameters u_a . We use random values to ensure that the results obtained are not because of a specific input structure. Travel time functions are given by (9) with $t_a = 19$, $c_a = 100$, and $\beta = 1$. We created three instances, called GridA, GridB, and GridC, by choosing the u_a 's from independent, identical, and uniform distributions for every arc $a \in A$ (0 to 11, 3 to 14, and 0 to 7, respectively). All these instances consider a demand of 100 users from the lower-left node to the upper-right node that form a single OD pair. For example, the picture on the left of Figure 2 depicts the topology and the variabilities of instance GridA. We computed the robust and added-variability equilibria corresponding to the three instances and the various values of Γ and ϕ . For the most part, we present the results for GridA, as results for the other two were similar. The complete set of inputs and outputs for all instances, including GridB and GridC, can be found in the online supplement to the paper (see footnote 8).

Figure 2 also presents the solutions corresponding to nominal and robust equilibria for some values of Γ . Under a nominal equilibrium, the flow is uniformly distributed across the grid because all arcs have the same expectation and this solution disregards travel-time uncertainty. The three solutions on the right correspond to robust equilibria for increasing values of Γ . The more attention users pay to uncertainty, the less they tend to use arcs with highly uncertain travel times. Table 2 complements these outputs by presenting summary statistics of travel times, which were computed using the simulation procedure previously described. It can be seen from the table that users tend to take longer paths instead of uncertain ones. This results in a higher mean but smaller variability. Indeed, we can see that for values of Γ around three, the variance tend to be smallest. The column “unfairness” represents the gap between the 95th and the 5th percentile of the travel times. The value for GridA when $\Gamma = 3$ is 1.226, down from 1.276 obtained for the nominal equilibrium. This represents a reduction of $1 - (1.226 - 1)/(1.276 - 1) \approx 18\%$. (The reduction is with respect to the lower bound, which equals one.) In addition, the standard deviation is reduced approximately 10.5%. The effect in GridC is less significant than in the others because there is much less variability in that instance.

Because all users have the same risk preference and all paths are eight arcs long, one can expect that added-variability equilibria approximate percentile equilibria. Figure 3 depict these equilibria, which resemble those in Figure 2 on the aggregate level. We will next further analyze both classes of equilibria.

Figure 4 shows the travel time distributions for the various equilibria. These distributions complement

Table 2 Summary Statistics of the Travel Time Distribution for Robust Equilibria

Γ	GridA				GridB				GridC			
	0	1	3	5	0	1	3	5	0	1	3	5
Mean	156.0	158.8	165.5	167.0	157.5	158.0	163.1	165.2	157.0	158.2	161.2	161.3
Stdev	11.5	11.3	10.3	10.0	10.6	10.0	9.8	11.4	4.85	4.35	4.43	4.42
Unfairness	1.276	1.259	1.226	1.223	1.247	1.223	1.214	1.251	1.108	1.096	1.098	1.098

Note. The column “unfairness” represents the variability of the middle 90% of observations, computed as 95th percentile / 5th percentile.




Figure 3 Added-Variability Equilibria of Instance GridA for $\phi = 0.25, 0.5, 0.75, 1$
 Note. The case of the nominal equilibrium coincides with $\Gamma = 0$ in Figure 2.




Figure 4 Cumulative Distribution Functions of Travel Time for the Robust (Left) and Added-Variability (Right) Equilibria of Instance GridA

the conclusion drawn from Table 2: when users are risk averse, the average travel time increases (curves are shifted to the right) but variability decreases (curves are steeper). For example, Figure 4 says that for the nominal equilibrium, around 50% of the time, a user will experience a travel time of 155 or less. For the robust equilibrium with $\Gamma = 3$, the corresponding probability is approximately 20%.

Because all these outputs are quite similar, which makes them hard to compare at the aggregate level, Figure 5 displays the distributions of regret for the solutions we computed. Each graph compares various solutions for a fixed percentile level. Notice that these graphs are piecewise constant because all flow taking the same path is indistinguishable and will experience the same regret. As expected, under a nominal equilibrium users have an almost constant 50th-percentile regret equal to 1.¹⁰ A priori there is no reason to expect that any of the other distributions shown in these graphs is vertical because robust and added-variability equilibria do not explicitly compute any percentiles. However, many of them show a low value

of regret for most users. The graphs provide evidence that more risk aversion calls for larger values of Γ or ϕ . Indeed, there is a strong correlation between the steepest distribution in a graph and the percentile level used to generate the graph. The conclusion is that appropriate parameters lead to a solution close to a percentile equilibrium. The proposed models handle uncertainty in a way that prevents risk-averse users from taking routes that can prove too long with high probability.

To capture the insights provided by Figure 5, Figure 6 presents the mean and the standard deviation of regrets for different percentile levels. The figure shows which robust model should be used for a certain risk profile. The percentile levels that correspond to the minima in these curves are increasing with Γ . For instance, one should use $\Gamma = 0$ for percentile levels less than 70th, $\Gamma = 1$ for levels between 70th and 95th, and $\Gamma = 3$ for levels higher than 95th. A similar observation can be made for added-variability equilibria.

6.2. Sioux Falls Instance

The instance Sioux Falls is a well-studied network in the transportation literature that is available at a repository managed by Bar-Gera (2002). The graph on the left of Figure 7 shows a representation of the Sioux Falls network. Because links are bidirectional, we represent each one with a pair of opposing arcs. Sioux

¹⁰ The distribution is not exactly vertical because the empirical 50th-percentile travel time of some users is slightly higher than the minimum possible. This deviation is attributable to sampling error arising from the Monte Carlo simulation.

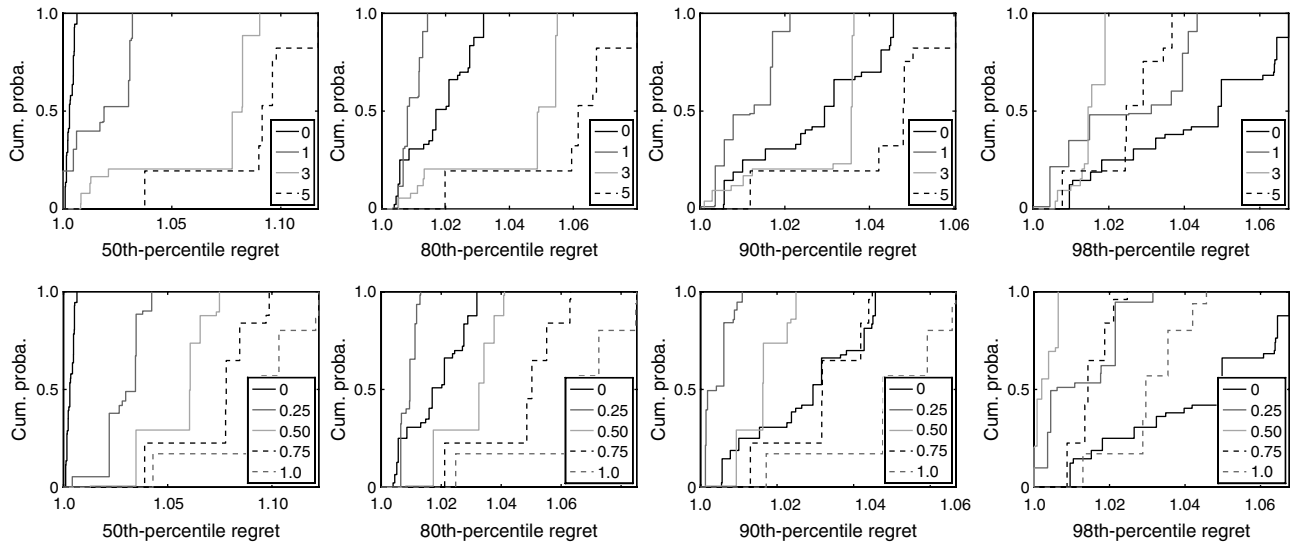


Figure 5 Distribution of Regret for Instance GridA

Note. *Top:* Robust equilibria. Each graph contains the curves for various values of Γ and a fixed percentile level. *Bottom:* Added-variability equilibria. Each graph contains the curves for various values of ϕ and a fixed percentile level.

Falls also uses BPR travel time functions. The original instance uses the standard $\beta = 0.15$; for the modified instance that we describe in the next section, we set $\beta = 1$ so the flow levels have a larger impact on travel times.

Because the original Sioux Falls instance does not contain any uncertainty parameters, we created artificial travel time uncertainty parameters u_a . The graph on the right of Figure 7 shows these parameters. Because longer arcs tend to have more uncertainty, we set u_a proportional to the free-flow travel time for most arcs. Typically the downtown area of a city is more congested as more cars use metered parking, enter and exit parking lots, look for parking places, there are more cabs that stop or circulate slowly, etc. We can thus argue that because of the larger number of random events that can affect traffic, travel times in downtown tend to be more uncertain than in the periphery. We arbitrarily labeled the Sioux Falls’

downtown as the vicinity of nodes 10, 16, and 17 and correspondingly assigned higher uncertainty values to arcs between those nodes. Similarly, we assigned medium uncertainty values to arcs in the proximity of that first group with the intention of modeling a transition area. The second group is composed of arcs between nodes 6, 5, 4, 11, 14, 15, and 19.

6.2.1. Sioux Falls with Simplified Demands. In this section we use a simplified version of Sioux Falls, referred to by Sfs, that we generated by considering five artificial OD pairs. The purpose of this simplification is to analyze the OD pairs one by one to develop insights for the trade-offs and benefits of robust equilibria. In §6.2.2, we present the results for the original Sioux Falls instance with all 528 OD pairs, referred to by SfC.

The five OD pairs we created are $\{(1, 19); (13, 8); (6, 21); (12, 18); (4, 17)\}$. We selected these five OD pairs so destinations are sufficiently far from

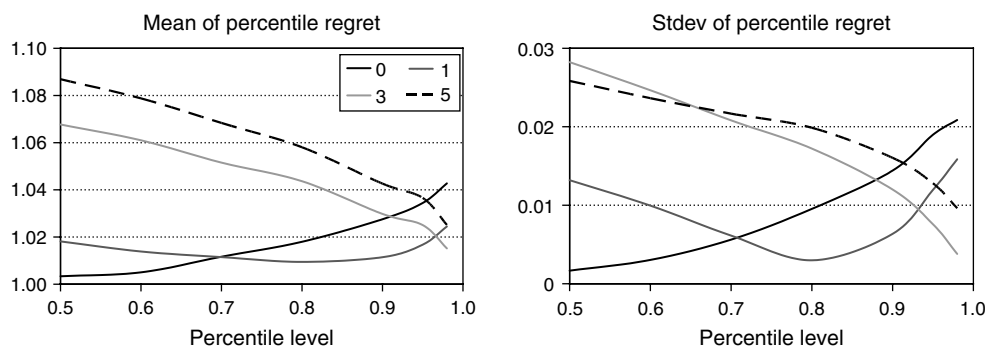


Figure 6 Mean and Standard Deviation of Percentile Regret for Robust Equilibria of Instance GridA

Note. Each curve represents a different value of Γ .




Figure 7 Sioux Falls Instance

Notes. Left: Network description. Labels represent node and link numbers. Right: Worst-case deviations u_a .

sources to ensure that, for large values of Γ , the robust equilibrium is not obtained by simply adding u_a to all arcs. Figure 8 shows the nominal and some of the robust equilibria; at the aggregate level, solutions for other values of Γ are similar so we omit them. Robust solutions tend to make more use of arcs in the periphery as they have low uncertainty whereas less flow is routed along arcs close to the downtown area. As an example take arc (10, 17) whose travel time is very uncertain. In the nominal solution, much of the flow with final destination 17 gets routed through it. Instead, the robust solution with $\Gamma = 1$ prefers routes along arc (16, 17), whereas those with larger Γ rather use the detour (19, 17) because (16, 17) also has a large uncertainty.

The graphs on the left of Figure 9 summarize these solutions by showing the cumulative probability distributions of total travel time for each OD pair. In those graphs, it can be seen how the risk aversion influences the equilibrium outcomes for the different OD pairs. Depending on the OD pair, robust solutions can have cumulative distributions that are similar to (second and third graphs in Figure 9 and the one not displayed by Figure 9), that dominate (fourth graph), or that are dominated by (first graph) those of the nominal equilibrium. As an example, let us explain the effect in the fourth graph. The corresponding OD pair goes to node 17, in downtown. Because arcs close to downtown have large u_a , the travel time of this OD pair has a large variance. For larger values of




Figure 8 Nominal and Robust Equilibria of Instance Sfs for $\Gamma \in \{1, 5\}$

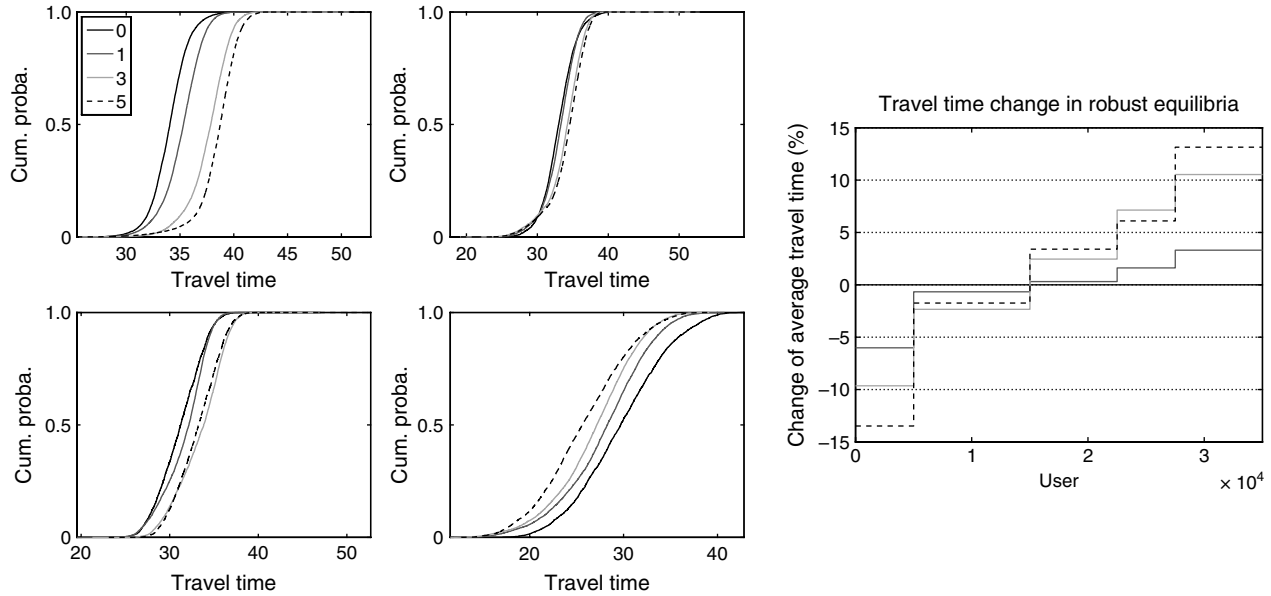


Figure 9 Travel Time Under Robust Equilibria of Instance SfS for Different Values of Γ .

Notes. Left: Cumulative distribution functions. Each graph depicts a different OD pair. One OD pair is not displayed because the curves corresponding to different values of Γ are indistinguishable and very similar to the second picture. Right: Change in expected travel times of robust equilibria with respect to those under a nominal equilibrium.

Γ the rerouting to arcs that are less variable achieves a smaller variability overall as evidenced in the distribution. Notice that one curve is to the right of the other does not say that the solution is or is not closer to an equilibrium; it just says that if travelers were controlled by a centralized coordinating agency, they would prefer that curve. Because users in a transportation network are not coordinated, we compare the realized travel time of a user to either the perceived utility she may have experienced had she acted differently (regret) or to the realized travel time of another user in the same situation (unfairness).

The graph on the right of Figure 9 presents the change between expected travel times in a robust solution and those in the nominal equilibrium. The graph is also piecewise constant because all users of the same OD pair are indistinguishable, making them experience the same distribution of travel times. Note that each curve is sorted independently to make them nondecreasing, and this means that users along each curve are not in one-to-one correspondence. For example, looking at the curve corresponding to $\Gamma = 5$, for some users there is a 14% reduction in the average travel time compared to the nominal equilibrium, for some there is small reduction of around 1%, and for the rest there is an increase of between 4% and 13%. For $\Gamma = 1$, the changes are milder, between -6% and 3% .

The top row of Figure 10 plots the distributions of regret for robust equilibria of instance SfS. As expected, we see that the steepest curve in the graph corresponding to the 50th percentile is that of the

nominal equilibrium. The robust equilibrium with $\Gamma = 1$ dominates the graph corresponding to the 90th percentile, and almost dominates for the other percentile levels (competing with $\Gamma = 0$ and 3 respectively). The conclusion, which agrees with the findings for the Grid instances, is that more risk aversion calls for using higher values of Γ .

The bottom row of Figure 10 shows similar graphs for added-variability equilibria. Looking again at the steepest curves in each graph, some of these solutions seem to be good approximations of percentile equilibria. As it happened for robust equilibria, more risk aversion calls for using higher values of ϕ . A difference between the solution concepts is that although for robust equilibria the solution with $\Gamma = 1$ seemed to be quite good for the different levels of risk aversion (80th, 90th, and 98th), the best solution among added-variability equilibria is more sensitive to the risk aversion of the users (but that of $\phi = 0.5$ is quite good).

6.2.2. Sioux Falls Complete. Now we turn to the original Sioux Fall instance with the original demands consisting of 528 OD pairs. The picture on the left of Figure 11 shows the robust equilibrium corresponding to $\Gamma = 1$. Robust equilibria for other values of Γ and added-variability equilibria are omitted because the aggregated flows look alike, despite differences in their path compositions. These differences are easier to visualize through other pictures. The uncertainty of travel times produces a similar effect on the flows as those noted for the simplified instance, although

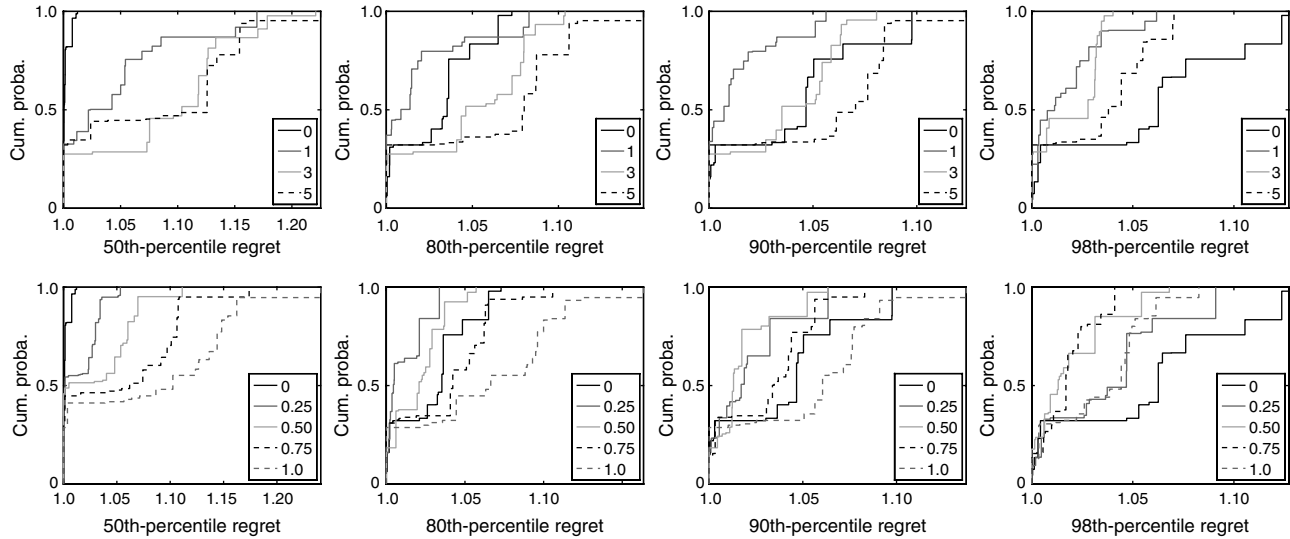


Figure 10 Distribution of Regret for Instance SfS

Note. Top: Robust equilibria. Each graph contains the curves for various values of Γ and a fixed percentile level. Bottom: Added-variability equilibria. Each graph contains the curves for various values of ϕ and a fixed percentile level.

these effects are more difficult to appreciate in the figure because of the interactions between all different OD pairs. We will not provide a user-by-user analysis of the empirical distribution as we did before because the quantity of OD pairs makes it prohibitive. Instead, we focus on aggregate measures and compare them across solutions.

The picture on the right of Figure 11 presents the change in expected travel time between a robust solution and the nominal one. (This graph is equivalent to that on the right of Figure 9.) This graph shows that slightly less than half of the users are better off and slightly more than half of the users incur in a longer expected travel time. These graphs suggest that robust solutions are somewhat similar with each other. The likely explanation is that nodes of many OD pairs are

close to each other, hence there are not many alternative paths that users can realistically consider. Hence, most of the effect of robust solutions is observed when increasing Γ from zero to one. This is confirmed by the distributions of percentile regret shown in the first row of Figure 12. Notice that the axis of the cumulative probabilities starts at 0.7 to facilitate the visualization of the curves. Indeed, more than 70% of the users are taking shortest paths with respect to any percentile of travel times and for any of the solutions we computed. For the rest of the users, robust equilibria balance the percentile regret corresponding to the risk tolerance of the users.

The bottom row of Figure 12 presents the regret distributions for added-variability equilibria. As for the simplified version, this method also approxi-

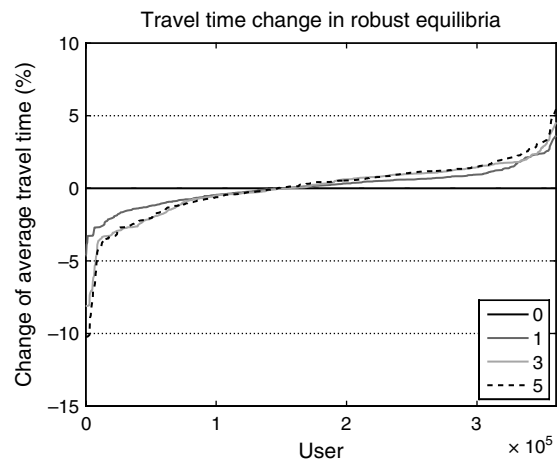



Figure 11 Solution to Instance SfC

Note. Left: Robust equilibrium for $\Gamma = 1$. Right: Change in expected travel times of robust equilibria with respect to those under a nominal equilibrium

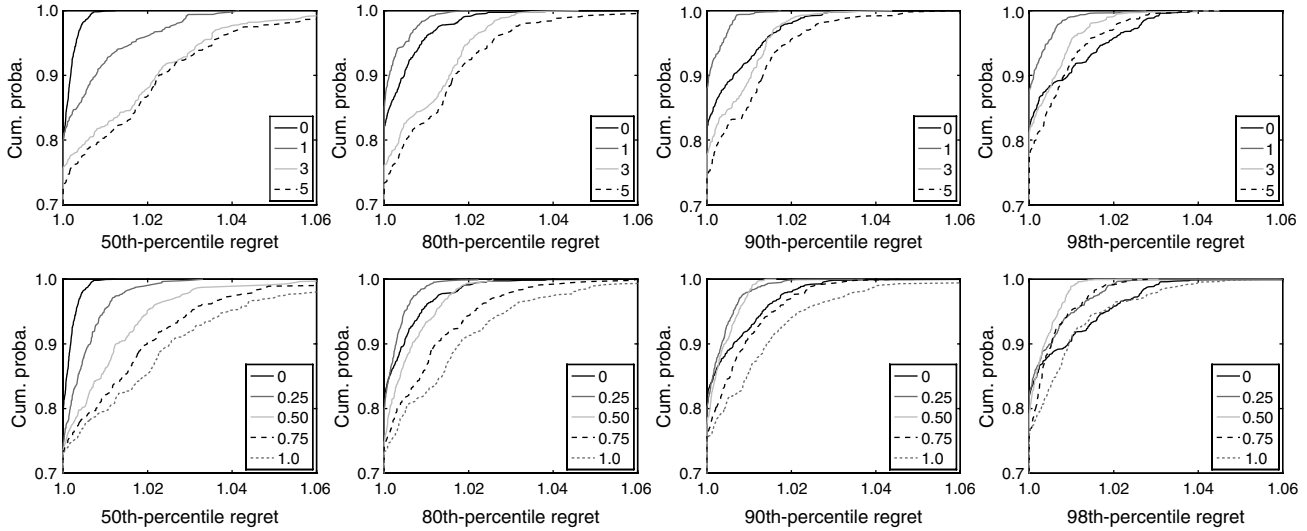


Figure 12 Distribution of Regret for Instance S1C

Notes. Top: Robust equilibria. Each graph contains the curves for various values of Γ and a fixed percentile level. Bottom: Added-variability equilibria. Each graph contains the curves for various values of ϕ and a fixed percentile level.

mates quite well a percentile equilibrium. Comparing both rows in the figure, we see again that although the robust equilibrium with $\Gamma = 1$ is the best approximation to percentile equilibrium for all percentile levels, the best added-variability equilibrium solution is somewhat sensitive to the percentile level.

6.3. Friedrichshain Instance

The most realistic instance we have considered, referred to as Fri, represents an area of Berlin called Friedrichshain. This network, shown on the left of Figure 13, was first presented in Jahn et al. (2005). Even though it is an order of magnitude larger than the previous ones, we can still solve the robust equilibrium problems without any difficulties albeit the exponential number of paths in the network. The network consists of 224 vertices, 523 arcs, and 506 different OD pairs. Free-flow travel times and capacities

were provided by Daimler Chrysler, and β in the BPR travel time functions was set to the standard value of 0.15. As the original data did not include uncertainty, we set u_a equal to 50% of the free-flow travel time for each arc $a \in A$ because longer arcs tend to have more variability. On the right of Figure 13, we show the robust equilibrium corresponding to $\Gamma = 1$, which is the one closer to being a percentile equilibrium, as shown below. We omit the other solutions because they are almost indistinguishable from this one at the aggregate level.

The top row of Figure 14 shows the distributions of regret for robust equilibria, whereas the one in the bottom shows those of added-variability equilibria. Notice that, as before, the vertical axes of these graphs start at 0.6 because approximately 70% of the users travel along shortest paths with respect to the corresponding percentile of travel times. As the figure shows, added-variability equilibria for this instance




Figure 13 Solution for Instance Fri

Note. Left: Worst-case deviations u_a . Right: Robust equilibrium for $\Gamma = 1$.




Figure 14 Distribution of Regret for Instance Fri

Note. Top: Robust equilibria. Each graph contains the curves for various values of Γ and a fixed percentile level. Bottom: Added-variability equilibria. Each graph contains the curves for various values of ϕ and a fixed percentile level.

are all dominated by the nominal equilibrium. Instead, the robust equilibrium corresponding to $\Gamma = 1$ dominates all the other curves for percentiles larger than 50. Hence, that solution is the best approximation to a percentile equilibrium for the case of risk-averse users.

Figure 15 shows that the robust equilibrium for $\Gamma = 1$ supports shorter travel times for approximately a third of the users and longer ones for approximately 20% of them. The magnitude of the changes is around 5% in different directions for both groups. Values of Γ equal to three or five provide solutions that are more at disequilibrium for all percentiles and that are more inefficient because half of the users travel significantly longer. The regret under those solutions,

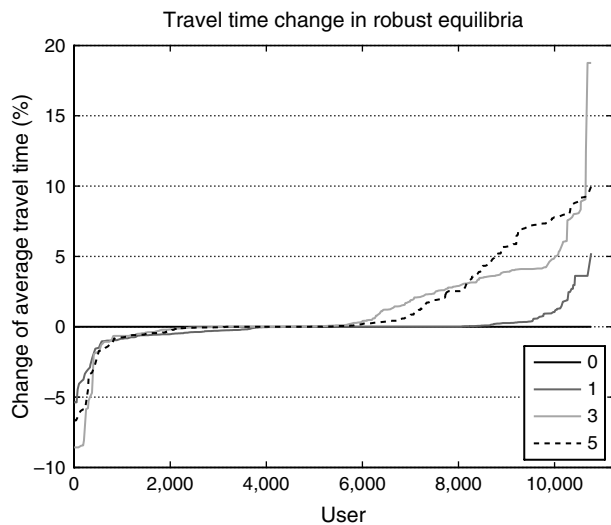


Figure 15 Change in Expected Travel Times of Robust Equilibria of Instance Fri with Respect to Those Under a Nominal Equilibrium

albeit better than that of a nominal equilibrium, is worse than in the case of $\Gamma = 1$.

7. Conclusions

Because the actual conditions of a networked system are subject to unforeseen events or influences, planners that want to predict user behavior should consider the existence of some degree of travel time uncertainty. The models of equilibrium under uncertainty that we have introduced allow the planner to incorporate some of the effects that this uncertainty may have on the route choices. We have shown through an extensive set of simulations that well-chosen robust and added-variability equilibria can approximately balance the duration of the trips with their variability, thereby diminishing the regret suffered by travelers. In all the instances considered in the computational study of §6, an optimal choice (or at least one of the best) when users are risk averse was to consider a robust equilibrium for $\Gamma = 1$. Larger values of Γ were not useful in improving the percentile regret for our set of instances because users become too pessimistic for the sizes of those networks and the distributions of travel times. It is likely that one needs to consider larger values of Γ for larger instances than those we have studied. We have provided some evidence that more risk aversion necessitates higher values of Γ when considering robust equilibria and higher values of ϕ when considering added-variability equilibria. Robust equilibria with $\Gamma = 1$ tend to provide solutions with low regret for multiple percentile levels whereas the solution with lowest regret among added-variability equilibria seems to be more sensitive to the percentile level chosen.

The computation of added-variability equilibria has the same complexity as computing an equilibrium that ignores uncertainty because it is a standard (deterministic) Wardrop equilibrium with respect to nominal travel times plus a constant fraction of u_a . Although computing robust equilibria is more difficult than that, we have provided evidence that they can be computed in practice for instances of moderate size using a column-generation algorithm. In fact, in our experience the bottleneck of our computations was generating the samples used in determining the regret and not computing the flows at equilibrium. We leave the efficient computation of percentile equilibria as an open problem.

Although in some cases travel times in an equilibrium under uncertainty are smaller than in a nominal equilibrium, we have seen that this is not true in general. When users are risk averse, they may feel compelled to select detours that are longer in expectation but that have smaller variability. The fact that travel times increase does not validate or invalidate the equilibrium assumptions. If anything, this may indicate that risk averseness may drive equilibria even further away from a system-optimal operating point than it was believed. Actually, recent research focuses in finding the extent of the inefficiency of equilibria in the deterministic Wardrop model (Roughgarden and Tardos 2002). Related to this, Bertsimas and Sim (2003) studied the tradeoff between the degree of robustness and efficiency. A topic of future research that relates to the model we have proposed is to study the relation between the degree of risk aversion and the quality of the resulting equilibrium. Finally, the network equilibrium model considered in this work assumes that there is an infinite number of users (i.e., the game is nonatomic) and that the demand is deterministic. Relaxing these assumptions is also left for future research.

Acknowledgments

The authors thank Michele Aghassi, whose constructive comments helped to improve this paper, and two anonymous referees for their careful reading of the manuscript and for numerous suggestions, including the example of traffic lights and the convex optimization formulation of (3). An extended abstract of this article appears in NET-COOP 2007 (Ordóñez and Stier-Moses 2007a).

Appendix A. Pricing Mechanism for Equilibria with Uncertain Travel Times

In 1952, William Vickrey, winner of the Nobel Prize for Economics, recommended that users of a transportation system pay tolls to achieve a more desired overall usage of the network. Indeed, a system planner can use tolls to make users internalize the congestion externalities they produce (Bergendorff, Hearn, and Ramana 1997). In this way, the resulting equilibrium will match the socially desirable solution because users will take the toll into account when

making their route decisions. It is well known that a fully efficient outcome can be achieved for Wardrop's user choice model by charging tolls equal to the marginal cost under a socially optimal solution. Independently, Yang and Huang (2004) and Fleischer, Jain, and Mahdian (2004) extended these ideas to the case of users with heterogeneous values of time. They characterize the aggregate flows that are enforceable using tolls, and offer a method to compute those tolls through linear programming duality (see Marcotte and Zhu (2009) for a recent treatment on this research).

We now discuss how the approach of Fleischer, Jain, and Mahdian (2004) can be used to show that tolls that induce optimal flows exist for the case of the solution concept introduced in §2. This approach also provides a method to compute these tolls. A given aggregate flow g is said to be *minimal* if the constraints (10b) are tight for every $a \in A$ under an optimal solution to the following linear program:

$$\min_h \sum_{P \in \mathcal{P}} l_P(g) h_P \quad (10a)$$

$$\text{s.t.} \quad \sum_{a \in P \in \mathcal{P}} h_P \leq g_a \quad a \in A \quad (10b)$$

$$\sum_{P \in \mathcal{P}_k} h_P = d_k \quad k \in K \quad (10c)$$

$$h_P \geq 0 \quad P \in \mathcal{P}. \quad (10d)$$

Fleischer, Jain, and Mahdian (2004) extend the previous formulation for the case of heterogeneous values of time and show that g is enforceable if and only if it is minimal (Theorem 3.1). The tolls necessary to obtain g turn out to be the optimal dual variables corresponding to (10b). This implies that tolls that lead to a socially optimal solution can be easily computed, extending what was already known for settings with a single value of time (Corollary 3.3).

The idea of using the dual variables corresponding to (10b) as tolls also works in the framework of uncertain travel times of the form $l_P(\cdot) + \delta_P$. The target flow depends on the model of uncertainty used by the system planner, which may be different from the users' because of different information availability. We consider two cases: (a) the planner has perfect information and thus aims to enforce the socially optimal flow, and (b) the planner has its own model of uncertainty for the whole network and aims for a robust social optimum flow. Regarding (a), a social optimum is defined as a solution to the following nonlinear program:

$$\min \sum_{P \in \mathcal{P}} l_P(h) h_P$$

$$\text{s.t.} \quad \sum_{P \in \mathcal{P}_k} h_P = d_k \quad k \in K$$

$$h_P \geq 0 \quad P \in \mathcal{P}.$$

The case (b) generalizes (a) by allowing the system planner to consider uncertain travel times, although not necessarily the same uncertainty as the users. Indeed, the system planner represents the travel time uncertainty vector η as nonnegative deviations from a nominal travel time that satisfy aggregate correlation conditions $\mathcal{U} \subseteq \mathbb{R}_{\geq 0}^m$. For example, \mathcal{U} can be defined as the set such that the total deviation

is bounded by Γ , in a similar fashion as we did for the robust equilibrium model (although we highlight that the total deviation considers all arcs in the network whereas in the robust equilibrium model it considered Γ arcs in each path). The robust social optimum is a solution that has minimal worst-case cost:

$$\begin{aligned} \min_h \quad & \max_{\eta} \left\{ \sum_{P \in \mathcal{P}} (l_P(h) + \sum_{a \in P} \eta_a) h_P : \eta \in \mathcal{U} \right\} \\ \text{s.t.} \quad & \sum_{P \in \mathcal{P}_k} h_P = d_k \quad k \in K \\ & h_P \geq 0 \quad P \in \mathcal{P}. \end{aligned}$$

Modifying Problem (10) by making the objective function use travel time functions $l_P(\cdot) + \delta_P$, it is possible to show that both the social optimum and the robust social optimum defined above are enforceable with tolls. One must show first that if these solutions do not turn out to be minimal, they could be reduced to make them minimal. Then, one can show that the optimal dual variables corresponding to (10b) provide the tolls that are needed to enforce these two solutions. We summarize the statements of these results in the propositions below.

PROPOSITION A.1. *The social optimum is enforceable with tolls when users follow the equilibrium model with uncertain travel times.*

PROPOSITION A.2. *The robust social optimum for an arbitrary uncertainty set \mathcal{U} is enforceable with tolls when users follow the equilibrium model with uncertain travel times.*

The tolls necessary to enforce a minimal flow h^* are given by the optimal dual variables to (10b) for the modified version of (10). The dual problem is as follows:

$$\max_{t, \pi} \quad \sum_{k \in K} \pi_k d_k - \sum_{a \in A} f_a^* t_a \quad (11a)$$

$$\text{s.t.} \quad \pi_k \leq l_P(h^*) + \delta_P + \sum_{a \in P} t_a \quad P \in \mathcal{P}_k, k \in K \quad (11b)$$

$$t_a \geq 0 \quad a \in A, \quad (11c)$$

where the dual variables are t_a for $a \in A$ and π_k for $k \in K$. This is a linear program with a large number of constraints, one for each possible path in \mathcal{P} . This problem can be solved with a constraint generation algorithm that considers a subset of paths at each iteration. Letting (t, π) be the optimal dual variables, one can solve a shortest path problem $SP(h)$ with travel time functions given by $l_P(h) + \delta_P + \sum_{a \in P} t_a$ to find if there is a path that violates any constraint (11b). If such a path exists it is added to the restricted master problem and the next iteration starts. Otherwise, when no violating path can be found, the current dual solution is optimal. We implemented this algorithm with AMPL and CPLEX and it obtained optimal taxes for all the instances that were described in §6.

References

Aashtiani, H. Z., T. L. Magnanti. 1981. Equilibria on a congested transportation network. *SIAM J. Algebraic Discrete Methods* 2(3) 213–226.

Aghassi, M., D. Bertsimas. 2006. Robust game theory. *Math. Programming* 107B 231–273.

Altman, E., T. Boulogne, R. El-Azouzi, T. Jiménez, L. Wynter. 2006. A survey on networking games in telecommunications. *Comp. Oper. Res.* 33(2) 286–311.

Andreatta, G., L. Romeo. 1988. Stochastic shortest paths with recourse. *Networks* 18 193–204.

Ashlagi, I., D. Monderer, M. Tennenholtz. 2006. Resource selection games with unknown number of players. *Proc. Fifth Internat. Joint Conf. Autonomous Agents Multiagent Systems*. ACM, New York, 819–825.

Atamtürk, A., M. Zhang. 2007. Two-stage robust network flow and design under demand uncertainty. *Oper. Res.* 55(4) 662–673.

Bar-Gera, H. 2002. Transportation network test problems. Accessed November 2009, <http://www.bgu.ac.il/~bargera/tntp/>.

Beckmann, M. J., C. B. McGuire, C. B. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.

Bell, M. G. H., C. Cassir. 2002. Risk-averse user equilibrium traffic assignment: An application of game theory. *Transportation Res.* 36B(8) 671–681.

Ben-Tal, A., A. Nemirovski. 1997. Robust truss topology design via semidefinite programming. *SIAM J. Optim.* 7(4) 991–1016.

Ben-Tal, A., A. Nemirovski. 1998. Robust convex optimization. *Math. Oper. Res.* 23(4) 769–805.

Ben-Tal, A., B. Golany, A. Nemirovski, J.-P. Vial. 2005. Supplier-retailer flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7(3) 248–273.

Bergendorff, P., D. W. Hearn, M. V. Ramana. 1997. Congestion toll pricing of traffic networks. P. M. Pardalos, D. W. Hearn, W. W. Hager, eds. *Network Optimization, Lecture Notes in Economics and Mathematical Systems*, Vol. 450. Springer, Berlin, 51–71.

Bertsekas, D. P., J. N. Tsitsiklis. 1991. An analysis of stochastic shortest path problems. *Math. Oper. Res.* 16(3) 580–595.

Bertsimas, D., M. Sim. 2003. Robust discrete optimization and network flows. *Math. Programming Ser. B* 98 49–71.

Bertsimas, D., A. Thiele. 2004. A robust optimization approach to supply chain management. D. Bienstock, G. Nemhauser, eds. *Proc. 10th Internat. Integer Programming Combin. Optim. Conf., IT Lecture Notes in Computer Science*, Vol. 3064. Springer, Berlin/Heidelberg, 86–100.

Bertsimas, D., A. Thiele. 2006. A robust optimization approach to inventory theory. *Oper. Res.* 54(1) 150–168.

Braess, D. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12 258–268. [English translation: Braess, D., A. Nagurney, T. Wakolinger 2005. On a paradox of traffic planning. *Transportation Sci.* 39(4) 446–450].

Bureau of Public Roads. 1964. Traffic assignment manual. U.S. Department of Commerce, Urban Planning Division, Washington, DC.

Dafermos, S. C. 1980. Traffic equilibrium and variational inequalities. *Transportation Sci.* 14(1) 42–54.

Daganzo, C. Y. Sheffi. 1977. On stochastic models of traffic assignment. *Transportation Sci.* 11(3) 253–274.

de Palma, A., Y. Nesterov. 1998. Optimization formulations and static equilibrium in congested transportation networks. CORE Discussion Paper 9861, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.

de Palma, A., N. Picard. 2005. Route choice decision under travel time uncertainty. *Transportation Res.* 39A(4) 295–324.

Dial, R. B. 1971. A probabilistic multi-path traffic assignment algorithm which obviates path enumeration. *Transportation Res.* 5(2) 83–111.

El-Ghaoui, L., H. Lebret. 1997. Robust solutions to least-square problems to uncertain data matrices. *SIAM J. Matrix Anal. Appl.* 18 1035–1064.

El-Ghaoui, L., M. Oks, F. Oustry. 2003. Worst-case value-at-risk and robust portfolio optimization: A conic programming approach. *Oper. Res.* 51(4) 543–556.

El-Ghaoui, L., F. Oustry, H. Lebret. 1998. Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* 9(1) 33–52.

- Fan, Y. Y., R. E. Kalaba, J. E. Moore. 2005a. Arriving on time. *J. Optim. Theory Appl.* **127**(3) 497–513.
- Fan, Y. Y., R. E. Kalaba, J. E. Moore. 2005b. Shortest paths in stochastic networks with correlated link costs. *Comput. Math. Appl.* **49** 1549–1564.
- Ferris, M. C., O. L. Mangasarian, J. S. Pang. 2001. *Complementarity: Applications, Algorithms and Extensions*, Vol. 50, *Applied Optimization*. Kluwer Academic, Dordrecht, The Netherlands.
- Fleischer, L., K. Jain, M. Mahdian. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. *Proc. 45th Annual IEEE Sympos. Foundations Comput. Sci.* IEEE Computer Society Press, Los Alamitos, CA, 277–285.
- Fourer, R., D. M. Gay, B. W. Kernighan. 2002. *AMPL: A Modeling Language for Mathematical Programming*. Brooks/Cole, Pacific Groove, CA.
- Gabriel, S., D. Bernstein. 1997. The traffic equilibrium problem with nonadditive path costs. *Transportation Sci.* **31** 337–348.
- Goldfarb, D., G. Iyengar. 2003. Robust portfolio selection problems. *Math. Oper. Res.* **28**(1) 1–38.
- Harsanyi, J. C. 1967. Games with incomplete information played by “Bayesian” players, I–III. Part I. The basic model. *Management Sci.* **14**(3) 159–182.
- Harsanyi, J. C. 1968. Games with incomplete information played by “Bayesian” players, I–III. Part II. Bayesian equilibrium points. *Management Sci.* **14**(5) 320–334.
- Hayashi, S., N. Yamashita, M. Fukushima. 2005. Robust Nash equilibria and second-order cone complementarity problems. *J. Nonlinear Convex Anal.* **6** 283–296.
- Holmström, B., R. B. Myerson. 1983. Efficient and durable decision rules with incomplete information. *Econometrica* **51** 1799–1820.
- ILOG. 2005. CPLEX. <http://www.ilog.com>.
- Jahn, O., R. H. Möhring, A. S. Schulz, N. E. Stier-Moses. 2005. System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* **53**(4) 600–616.
- Liu, H., X. Ban, B. Ran, P. Mirchandani. 2002. An analytical dynamic traffic assignment model with stochastic network and travelers’ perceptions. *Transportation Res. Record* **1783** 125–133.
- Lo, H. K., Y.-K. Tung. 2003. Network with degradable links: Capacity analysis and design. *Transportation Res.* **37B**(4) 345–363.
- Marcotte, P., D. L. Zhu. 2009. Existence and computation of optimal tolls in multiclass network equilibrium problems. *Oper. Res. Lett.* **37** 211–214.
- Mirchandani, P. B., H. Soroush. 1987. Generalized traffic equilibrium with probabilistic travel times and perceptions. *Transportation Sci.* **21** 133–152.
- Nagurney, A. 1999. *Network Economics: A Variational Inequality Approach*, 2nd ed. Kluwer Academic, Dordrecht, The Netherlands.
- Nie, Y., X. Wu. 2009. Shortest path problem considering on-time arrival probability. *Transportation Res.* **43B** 597–613.
- Nikolova, E., M. Brand, D. R. Karger. 2006. Optimal route planning under uncertainty. *Proc. Internat. Conf. Automated Planning Scheduling*, Cumbria, UK, 131–141.
- Noland, R. B., K. Small, P. Koskenoja, X. Chu. 1998. Simulating travel reliability. *Regional Sci. Urban Econom.* **28**(5) 535–564.
- Ordóñez, F., N. E. Stier-Moses. 2007a. Robust wardrop equilibrium. *Proc. First Inter. Conf. Network Control Optim. Lecture Notes in Comput. Sci.*, Vol. 4465. Springer, Berlin, 247–256.
- Ordóñez, F., J. Zhao. 2007b. Robust capacity expansion of network flows. *Networks* **50**(2) 136–145.
- Patriksson, M. 1994. *The Traffic Assignment Problem: Models and Methods*. VSP, Utrecht, The Netherlands.
- Roughgarden, T. 2002. How unfair is optimal routing? *Proc. 13th Annual ACM-SIAM Sympos. Discrete Algorithms*, SIAM, Philadelphia, 203–204.
- Roughgarden, T., É. Tardos. 2002. How bad is selfish routing? *J. ACM* **49** 236–259.
- Sheffi, Y. 1985. *Urban Transportation Networks*. Prentice-Hall, Englewood, NJ.
- Smith, M. J. 1979. The existence, uniqueness and stability of traffic equilibria. *Transportation Res.* **13B** 295–304.
- Uchida, T., Y. Iida. 1993. Risk assignment: A new traffic assignment model considering risk of travel time variation. C. F. Daganzo, ed. *Proc. 12th Internat. Sympos. Transportation Traffic Theory*. Elsevier, Amsterdam, 89–105.
- Ukkusuri, S. V., S. T. Waller. 2010. Approximate expressions for network performance under demand uncertainty. *Transportation Lett.* Forthcoming.
- Vanderbei, R. J. 1999. LOQO: An interior point code for quadratic programming. *Optim. Methods Software* **12** 451–484.
- Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Institute of Civil Engineers, Part II*, Vol. 1, 325–378.
- Yang, H., H.-J. Huang. 2004. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Res.* **38B** 1–15.