Stochastic Selfish Routing

E. Nikolova[1] and N. E. Stier-Moses[2]

[1] MIT CSAIL, Cambridge, MA, USA, nikolova@mit.edu
[2] Columbia Business School, New York, NY, USA, stier@gsb.columbia.edu

Abstract. We embark on an agenda to investigate how stochastic delays and risk aversion transform traditional models of routing games and the corresponding equilibrium concepts. Moving from deterministic to stochastic delays with risk-averse players introduces nonconvexities that make the network game more difficult to analyze even if one assumes that the variability of delays is exogenous. (For example, even computing players' best responses has an unknown complexity [24].) This paper focuses on equilibrium existence and characterization in the different settings of atomic vs. nonatomic players and exogenous vs. endogenous factors causing the variability of edge delays. We also show that succinct representations of equilibria always exist even though the game is non-additive, i.e., the cost along a path is *not* a sum of costs over edges of the path as is typically assumed in selfish routing problems. Finally, we investigate the inefficiencies resulting from the stochastic nature of delays. We prove that under exogenous stochastic delays, the price of anarchy is exactly the same as in the corresponding game with deterministic delays. This implies that the stochastic delays and players' risk aversion do not further degrade a system in the worst-case more than the selfishness of players.

Keywords: Non-additive nonatomic congestion game, stochastic Nash equilibrium, stochastic Wardrop equilibrium, risk aversion.

1 Introduction

Heavy traffic and the uncertainty of traffic conditions exacerbate the daily lives of millions of people across the globe. According to the 2010 Urban Mobility Report [32], "in 2009, congestion caused urban Americans to travel 4.8 billion hours more and to purchase an extra 3.9 billion gallons of fuel for a congestion cost of $115 billion." For a comparison, that congestion cost was $85 billion in 1999. High and variable congestion necessitates drivers to *buffer in extra time* when planning important trips. The recommendation in the report was to consider a buffer of approximately 30% (Los Angeles) to 40% (Chicago) more than the average travel time, and around twice as long as the travel time at night when traffic is light.

A common driver reaction in the face of heavy and uncertain traffic conditions is to look for alternate, sometimes longer but less crowded and less variable routes [15]. With the widespread use of ever-improving technologies for measuring traffic, one might ask: is there a way to game the system? What route should be selected given other drivers' route choices? Considering routing games on networks where delay functions are stochastic, we analyze the resulting equilibria when strategic, risk-averse commuters take into account the *variability* of delays. This approach generalizes the traditional model of Wardrop competition [37] by incorporating uncertainty.

Risk aversion forces players to go beyond considering expected delays. Since it is unlikely that they base their routing decisions on something as complicated as a full distribution of delays along an exponential number of possible paths, it is reasonable that considering expected delays and their standard deviations is a good first-order approximation on route selection. To incorporate the standard deviation of delays into the players' objectives, we consider the traditional *mean-standard deviation* (mean-stdev) objective [14, 18] whereby players minimize the cost on a path, defined as the path mean plus a risk-aversion factor times the path standard deviation.[3] By linearity of expectations, the mean of the path equals the sum of the means over all its edges. However, the standard deviation along a path does not decompose as a sum over edges because of the *risk-diversification effect*. Instead, it is given by the square root of the sum of squared standard deviations on the edges of that path. Due to the complicating square root, a single player's subproblem—a shortest path problem with respect to stochastic costs—is a nonconvex optimization problem for which no polynomial running-time algorithms are known. This is in sharp contrast to the subproblem of the Wardrop network game—a shortest path problem, which admits efficient solutions.

A compelling interpretation of this objective in the case of normally-distributed uncertainty is that the mean-stdev of a path equals a percentile of delay along it. This model is also related to typical quantifications of risk, most notably the value-at-risk objective commonly used in finance, whereby one seeks to minimize commute time subject to arriving on time to a destination with at least, say, 95% chance.

Our mean-stdev model works for *arbitrary* distributions with finite first and second moment. To simplify the analysis, throughout this paper we assume that delays of different edges are uncorrelated. Nevertheless, a limited amount of correlation is to be expected in practice; for example, if there is an accident in a location, it causes ripple effects upstream. We remark that local correlations can be addressed with a polynomial graph transformation that encodes correlation explicitly in edges by modifying the standard-deviation functions with correlation coefficients [23]. This results in a graph with independent edge delays where all our results and algorithms carry through.

Related Work Our model is based on the traditional competitive network game introduced by Wardrop in the 1950's where he postulated that the prevailing traffic conditions can be determined from the assumption that players jointly select shortest routes [37]. The game was formalized in an influential book by Beckmann *et al.* that lays out the mathematical foundations to analyze competitive networks [4]. These models find applications in various domains such as transportation [33] and telecommunication net-

[3] Another alternative would have been to consider the mean-variance objective. This approach reduces to a deterministic Wardrop network game in which the edge delay functions already incorporate the information on variability. However, the mean and variance are measured in different units so a combination of them is hard to interpret. In addition, under this objective it may happen that players select routes that are stochastically dominated by others. Although this counterintuitive phenomenon may also happen under the mean-stdev objective with some artificially constructed distributions, it is guaranteed *not* to happen under normal distributions.

works [1]. In the last decade, these games received renewed attention with many studies aimed at understanding under what conditions equilibria exist, what uniqueness properties they satisfy, how to compute them efficiently, how expensive they are in relation to a centralized solution, and how to align incentives so the equilibria become optimal. For general references on these topics, we refer the readers to some recent surveys [9, 27].

In the majority of models used by theoreticians who study the properties of network games, and by practitioners who compute solutions to real problems, delays have been considered deterministic. Although there are models that incorporate some form of uncertainty [3, 5, 16, 17, 19, 36], none of these models has become widely accepted in practice, nor have they been extensively studied. Perhaps the only exception is the *stochastic user equilibrium* model, introduced by Dial in the 1970's [11], which has been studied and used in practice (see, e.g., [34, 35]). Under it, different players *perceive* each route differently, distributing demand in the network according to a logit model. To reduce route enumeration, the model just takes into account a subset of "efficient routes." In contrast, the objective of the players in the network game we consider is to choose the path that minimizes the mean plus a multiple of the standard deviation of delay. This problem belongs to the class of stochastic shortest path problems (see, for instance, some classic references [2, 6] and some newer ones [12, 13, 25, 22]).

In the network games literature, the model most related to our work is that of Ordóñez and Stier-Moses [28]. They introduce a game with uncertainty elements and risk-averse users and study how the solutions provided by it can be approximated numerically by an efficient column-generation method that is based on robust optimization. The main conclusion is that the solutions computed using their approach are good approximations of *percentile equilibria* in practice. Here, a percentile equilibrium is a solution in which percentiles of delays along flow-bearing paths are minimal. The main difference between their approach and ours is that their insights are based on computational experiments whereas the current work focuses on theoretical analysis and also considers the more general settings of endogenously-determined standard deviations and atomic games.

Next, we formally define our model and equilibrium concepts and study the existence of equilibrium under exogenous (Section 3) and endogenous (Section 4) variability of delays. We summarize these results in Table 1. We then prove that equilibria that use polynomially-many paths—referred to as succint—exist (Section 5), and finally we analyze properties of the socially-optimal solution and study the inefficiency of stochastic Wardrop equilibria (Section 6).

2 The Model

We consider a directed graph $G = (V, E)$ with an aggregate demand of d_k units of flow between source-destination pairs (s_k, t_k) for $k \in K$. We let $\mathcal{P}_k$ be the set of all paths between s_k and t_k, and $\mathcal{P} := \cup_{k\in K}\mathcal{P}_k$ be the set of all paths. We encode players decisions as a flow vector $\mathbf{f} = (f_\pi)_{\pi\in\mathcal{P}} \in \mathbb{R}_+^{|\mathcal{P}|}$ over all paths. Such a flow is feasible when demands are satisfied, as given by constraints $\sum_{\pi\in\mathcal{P}_k} f_\pi = d_k$ for all $k \in K$. For simplicity, when we write the flow on an edge f_e depending on the full flow $\mathbf{f}$, we

	Exogenous Noise	Endogenous Noise
Nonatomic Users	Equilibrium exists; Solves exponentially-large convex program	Equilibrium exists; Solves variational inequality
Atomic Users	Equilibrium exists; Potential game	No pure strategy equilibrium

Table 1. Equilibria in Stochastic Routing Games

refer to $\sum_{\pi \ni e} f_\pi$. When we need multiple flow variables, we use the analogous notation $\mathbf{x}, x_\pi, x_e$.

The congestible network is modeled with stochastic delay functions $\ell_e(x_e)+\xi_e(x_e)$ for each edge $e \in E$. Here, $\ell_e(x_e)$ measures the expected delay when the edge has flow x_e, and the random variable $\xi_e(x_e)$ represents the stochastic delay error. The function $\ell_e(\cdot)$ is assumed continuous and non-decreasing. The expectation of $\xi_e(x_e)$ is zero and its standard deviation is $\sigma_e(x_e)$, for a continuous and non-decreasing function $\sigma_e(\cdot)$. Although the distribution may depend on x_e, we will separately consider the simplified case in which $\sigma_e(x_e) = \sigma_e$ is a constant given exogenously, independent from x_e. We also assume that these random variables are all uncorrelated with each other. Risk-averse players choose paths according to the mean-standard deviation (mean-stdev) objective, which we also refer to as the cost along route π:

$$Q_\pi(\mathbf{f}) := \sum_{e\in\pi} \ell_e(f_e) + \gamma \sqrt{\sum_{e\in\pi} \sigma_e(f_e)^2}\,, \tag{1}$$

where $\gamma \geq 0$ quantifies the risk aversion of players, assumed homogeneous.

The *nonatomic* version of the game considers the setting where infinite players control an insignificant amount of flow each so the path choice of a player does not unilaterally affect the costs experienced by other players (even though the joint actions of players affect other players). The following definition captures that at equilibrium players route flow along paths with minimum cost $Q_\pi(\cdot)$.

Definition 1. *The* stochastic Wardrop equilibrium *of a nonatomic routing game is a flow* $\mathbf{f}$ *such that for every source-destination pair* $k \in K$ *and for every path* $\pi \in \mathcal{P}_k$ *with positive flow,* $Q_\pi(\mathbf{f}) \leq Q_{\pi'}(\mathbf{f})$ *for every path* $\pi' \in \mathcal{P}_k$ *.*

Instead, the *atomic* version of the game assumes that each player wishes to route one unit of flow. Consequently, the path choice of even one player directly affects the costs experienced by others. There are two versions of the atomic game: in the splittable case players can split their demands along multiple paths, and in the unsplittable case they are forced to choose a single path. In this paper we focus on the *atomic unsplittable* case, which we will sometimes refer to just as *atomic*. The natural extension of Wardrop equilibrium to the atomic case only differs in that players need to anticipate the effect of a player changing to another path. This game always admits a mixed-strategy equilibrium (under the standard expected payoffs with respect to the mixing probabilities) because it is a finite normal-form game [21], so we focus on the existence of pure-strategy equilibria.

Definition 2. *A pure-strategy* stochastic Nash equilibrium *of the atomic unsplittable routing game is a flow* $\mathbf{f}$ *such that for every source-destination pair* $k \in K$ *and for every path* $\pi \in \mathcal{P}_k$ *with positive flow, we have that* $Q_\pi(\mathbf{f}) \leq Q_{\pi'}(\mathbf{f} + \mathcal{I}_{\pi'} - \mathcal{I}_\pi)$ *for every* $\pi' \in \mathcal{P}_k$*. Here,* $\mathcal{I}_\pi$ *denotes a vector that contains a one for path* π *and zeros otherwise.*

One of the goals of this work is to evaluate the performance of equilibria. Hence, we define a social cost function that will allow us to compare different flows and determine the inefficiency of solutions. The social cost function is the total cost among players:

$$C(\mathbf{f}) := \sum_{\pi \in \mathcal{P}} f_\pi Q_\pi(\mathbf{f})\,. \tag{2}$$

3 Exogenous Standard Deviations

In this section, we consider exogenous noise factors, which result in constant standard deviations $\sigma_e(x_e) = \sigma_e$ that do not depend on the flow on the edge. In this case, the path cost (1) can be written as $Q_\pi(\mathbf{f}) = \sum_{e\in\pi} \ell_e(f_e) + \gamma(\sum_{e\in\pi} \sigma_e^2)^{1/2}$. We investigate the existence of equilibria and provide a characterization. First, we show that an equilibrium always exists, despite the challenge posed by the non-additive cost function. Due to space restrictions, missing proofs can be found in the full version of this paper.

Theorem 1. *A nonatomic routing game with exogenous standard deviations always admits a stochastic Wardrop equilibrium.*

The proof uses a path-based convex programming formulation given by Ordóñez and Stier-Moses [28]:

$$\min\Big\{\sum_{e\in E}\int_0^{x_e}\ell_e(z)dz + \sum_{\pi\in\mathcal{P}}\gamma f_\pi\sqrt{\sum_{e\in\pi}\sigma_e^2} : \text{such that } x_e = \sum_{\pi\in\mathcal{P}:\, e\in\pi} f_\pi \text{ for } e\in E,$$
$$d_k = \sum_{\pi\in\mathcal{P}_k} f_\pi \text{ for } k\in K,\ f_\pi \geq 0 \text{ for } \pi\in\mathcal{P}\Big\}. \tag{3}$$

Besides proving existence, the formulation also provides a way to compute this equilibrium using column generation. We remark that this method typically will not use many paths and hence, it is likely to be practical. In addition, the formulation implies that the equilibrium is unique, provided that the objective function (3) is strictly convex:

Corollary 1. *The equilibrium of the stochastic nonatomic routing game with exogenous standard deviations is unique (in terms of edge loads) whenever the expected delay functions are strictly increasing.*

We now return briefly to the question of computation. The convex program (3) contains exponentially-many variables (the flows on all paths) and a polynomial number of constraints. We will see in Section 5 that an equilibrium always has a succinct decomposition that uses at most $|E|$ paths; unfortunately, since we do not know ahead of time which paths these are, we cannot write a succinct version of the convex program. In the case of constant expected delays, the objective (3) coincides with the social cost, and both problems reduce to computing a stochastic shortest path for each source-destination pair. Thus, both the equilibrium and social optimum computation are at least as hard as the stochastic shortest path problem [26, 24].

Theorem 2. *When the expected delays and standard deviations are constant for each edge, the equilibrium and social optimum coincide and can be found in time* $n^{O(\log n)}$.

Now, we switch to the atomic unsplittable case and show that the stochastic routing game admits a potential function. We prove this using the characterization given by Monderer and Shapley [20]. The potential game structure implies that an equilibrium always exists.

Theorem 3. *An atomic unsplittable routing game with exogenous standard deviations is potential and, therefore, it always admits a pure-strategy stochastic Nash equilibrium.*

In contrast to the uniqueness of equilibrium in the nonatomic game, the pure-strategy equilibria in the atomic case need not be unique because they are a generalization of those in deterministic games, which admit multiple equilibria. (For example, a game with two players with a unit demand choosing among three parallel edges with $\ell_e(x_e) = x_e$ and $\sigma_e = 0$ admits three equilibria.)

4 Endogenous Standard Deviations

In this section, we consider flow-dependent standard deviations of edge delays. This makes the standard deviations endogenous to the game. We show that equilibria exist in the nonatomic game but they may not exist in the atomic game.

The following example illustrates how an equilibrium changes under endogenous standard deviations. Assume a demand of $d = 1$ and consider a network consisting of two parallel edges with delays $\ell_1(x) = x$ and $\ell_2(x) = 1$, followed by a chain of k edges that players must traverse. This instance admits two paths, each comprising one of the two parallel edges and the chain. We let L denote the expected delay along the chain (a constant since the flow traversing it is fixed) and assume that $\sigma_e(x_e) = sx_e$ for all edges, for some constant $s \geq 0$. Although both the deterministic and the exogenous standard-deviation games are equivalent to Pigou's instance [30, 31], the equilibrium with endogenous standard deviations changes significantly. Indeed, it is given by a root of the degree-4 polynomial $(1-4s^2)x^4+4s^2x^3+(4s^4-2s^2-4ks^2)x^2-4s^4x + s^4$ such that $x \in [0, 1]$, which in principle might not exist (where x denotes the flow on one of the two paths). An insight with algorithmic implications arising from this example is that an equilibrium in the stochastic game does not decompose to equilibria in subgraphs of the given graph, and in fact it may be quite different from the equilibria in the subgraphs. Hence, it is not immediate how to decompose the problem by partitioning a graph into smaller pieces.

We can show the existence of equilibrium in the general nonatomic setting via a variational inequality. In fact, the following results also hold in the much more general setting where edge-delay functions depend on the full vector of flows, as long as this dependence is continuous.

Theorem 4. *The nonatomic game with endogenous standard deviations admits a stochastic Wardrop equilibrium.*

In contrast to the case of exogenous standard deviations, however, the game with endogenous standard deviations is not *potential* [20] and equilibria cannot be characterized as the solution to a global optimization problem.

Proposition 1. *The stochastic routing game with nonconstant variances does not admit a cardinal potential.*

As in the deterministic case, the stochastic game may have multiple edge-flow equilibria when the delays and variances are not strictly increasing. An open question that remains is whether the equilibrium is unique when the expected delay and/or standard-deviation functions are strictly increasing. The standard approach to establish the uniqueness of a solution to a variational inequality is to show the monotonicity of the path-cost operator (the vector of path cost functions for all paths). However, neither monotonicity, nor a weaker notion of pseudo-monotonicity holds for our problem.

Proposition 2. *The path-cost operator of the nonatomic routing game with endogenous standard deviations is not pseudo-monotone.*

Although we were not able to prove uniqueness in general, we can do so in the extreme cases of players' risk attitudes. We do so by showing that in those cases the stochastic game resembles a deterministic one.

Proposition 3. *In the two extreme settings where players are either risk-neutral or infinitely risk averse, the nonatomic routing game admits a unique stochastic Wardrop equilibrium, for strictly increasing expected edge delays and standard deviations.*

In contrast to the existence of equilibria in nonatomic games, there are atomic games that do not admit a pure-strategy Nash equilibrium. Mixed-strategy equilibria always exist because the game is finite [21].

Proposition 4. *The atomic unsplittable routing game with endogenous standard deviations may not have pure-strategy Nash equilibria, even in the case of a single source-destination pair in a series-parallel network with affine edge mean and standard-deviation functions.*

5 Succinct Representations of Equilibria

We now turn our attention to how one can decompose an equilibrium of the nonatomic game represented as an edge-flow vector to a path-flow vector, and to whether a succinct vector of path flows at equilibrium always exists. The first question is trivial in the deterministic routing game: any path-flow decomposition of an edge flow at equilibrium works since path costs are additive. Instead, path costs of the stochastic game are non-additive and different flow decompositions of the same edge flow may incur different path costs. In particular, for an edge flow at equilibrium, some path-flow decompositions are at equilibrium and others are not. This is captured by the next lemma which illustrates that shortest paths with respect to path costs do not need to satisfy Bellman equations since a subpath of a shortest path need not be shortest.

Lemma 1. *In a nonatomic game, not all path-flow decompositions of an edge flow at equilibrium constitute an equilibrium.*

The previous lemma prompts the question of how one can find a flow decomposition of an equilibrium given as an edge flow. Does a succinct decomposition always exist (namely one that assigns positive flow to only polynomially-many paths)? The next few results provide positive answers to these questions. We show that succinct flow decompositions exist and they can be found in time slightly larger than polynomial, $|V|^{O(\log|V|)}$, which is the best-known running time of an exact algorithm for solving the underlying stochastic shortest path problem [26]. Alternatively, using a fully-polynomial approximation algorithm for the stochastic shortest path problem [24], one can find approximate flow decompositions of equilibria in polynomial time. We first provide a characterization of flow decompositions of equilibria that will enable us to show the existence of succinct decompositions.

Lemma 2. *Consider a flow decomposition $\mathbf{f}^P$ of an edge-flow equilibrium with support $P \subset \mathcal{P}$ (the set of paths with positive flow). Then, every flow decomposition whose support is a subset of P and whose resulting edge flow is the same as that of $\mathbf{f}^P$ is also at equilibrium.*

Using the lemma above, we can prove the existence of succinct equilibrium decompositions.

Theorem 5. *For an equilibrium $(f_e)_{e\in E}$ given as an edge flow, there exists a* succinct *flow decomposition that uses at most $|E|$ paths. Furthermore, this decomposition can be found in time $|V|^{O(\log|V|)}$, and an ϵ-approximate equilibrium succinct decomposition can be found in polynomial time.*

It remains open whether finding an equilibrium can be done in polynomial time. This is related to the open question of whether the stochastic shortest path subproblem is in P [26].

Corollary 2. *Given an edge flow, we can verify that it is at equilibrium in time $|V|^{O(\log|V|)}$.*

Analogously, we can verify that a given set of player strategies (paths) in the atomic setting forms an equilibrium in time $|V|^{O(\log|V|)}$.

6 Price of Anarchy

In this section we compute bounds for the price of anarchy (POA) for stochastic Wardrop equilibria of the nonatomic game. Recall that the price of anarchy is defined as the supremum over all problem instances of the ratio of the equilibrium cost to the social optimum cost [29]. In the case of exogenous standard deviations, the POA turns out to be the same as in the corresponding deterministic game: it is $4/3$ for linear expected delays and $(1-\beta)^{-1}$ for general expected delays for an appropriate definition of β as in Correa *et al.* [7] for the corresponding deterministic routing game. The bounds result from a modification of the bounding techniques of Correa *et al.* [7, 8].

In the case of endogenous standard deviations, an analysis of the price of anarchy is more elusive and it remains open whether the equilibrium is unique. For this reason, we focus our analysis to the limiting case of extreme risk aversion (the other extreme case, where users are risk neutral, is well-understood). Hence, we assume that path costs are equal to the path standard deviations $Q_\pi(\mathbf{f}) = (\sum_{e\in\pi} \sigma_e(f_e)^2)^{1/2}$. Recall that in this extreme case, Proposition 3 implies that there is a unique equilibrium that can be computed efficiently with a convex program.

We now show that the first order optimality conditions of the optimization problem that defines socially-optimal solutions are satisfied at the equilibrium, when standard-deviation delay functions are monomials of the same degree. Note that in the deterministic case, it is known that the POA is exactly one precisely for that class of delay functions [10].

Theorem 6. *Consider a nonatomic network game with endogenous standard deviations of the form $\sigma_e(x_e) = a_e x_e^p$ for some fixed $p \geq 0$. An equilibrium of the game is a stationary point of the social-optimum (SO) problem*

$$\min\Big\{ \sum_{\pi\in\mathcal{P}} f_\pi Q_\pi(\mathbf{f}) : \sum_{\pi\in\mathcal{P}_k} f_\pi = d_k \ \forall k\in K \text{ and } f_\pi \geq 0 \ \forall \pi\in\mathcal{P}\Big\}.$$

As a corollary from the above theorem, whenever the SO problem has a unique stationary point, it would follow that equilibria and social optima coincide and, consequently, the price of anarchy would be 1. Before we identify settings for which convexity of the social cost holds, we show that despite the somewhat misleading square root, the path costs are convex in the edge-flow variables when the standard deviations $\sigma_e(x_e)$ are convex functions.

Proposition 5. *The path costs $Q_\pi(\mathbf{x}) = (\sum_{e\in\pi} \sigma_e(x_e)^2)^{1/2}$ are convex whenever the edge standard-deviation functions $\sigma_e(x_e)$ are convex.*

As a corollary from Proposition 5 and from the fact that the sum of convex functions is convex, it follows that path costs are also convex in the general case where the path cost is a sum of the mean and standard deviation of the path, as long as the edge mean and standard-deviation functions are convex in the edge flow.

Next, we identify sufficient conditions for the convexity of the social cost, which bear an intriguing resemblance to the sufficient conditions for the uniqueness of equilibrium mentioned earlier.

Proposition 6. *The social cost $C(\mathbf{x}) = \mathbf{x}Q(\mathbf{x}) = \sum_{\pi\in\mathcal{P}} x_\pi Q_\pi(\mathbf{x})$ is convex whenever the path-cost operator Q is monotone and the path costs $Q_\pi(\mathbf{x})$ are convex.*

As established above, the path costs are convex (under convex standard-deviation functions), however the path-cost operator is not necessarily monotone even in the basic case of linear standard deviation functions equal to $\sigma_e(x) = x$, and thus, the social cost may not be convex as shown in Figure 1. Nevertheless, we can still show that the POA is 1 in a network of n pairs of parallel edges connected in series.

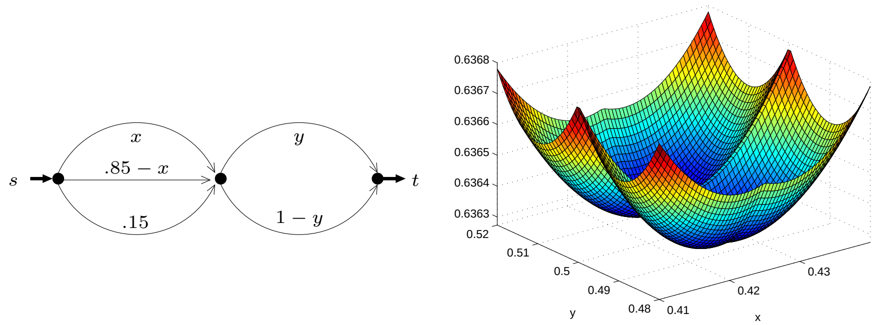

(a) Parametrization of edge flows (b) Social cost as a function of edge flows x and y

Fig. 1. Non-convex slice of the social cost function

Proposition 7. *Consider a nonatomic game on a network of n pairs of parallel edges connected in series with zero mean delays and standard deviation functions equal to $\sigma_e(x) = x$ for all edges. In this case, socially-optimal flows and equilibria coincide.*

Despite the limitation of the hypothesis, the proof requires a careful analysis to bound the social cost, in contrast to our results for exogenous standard deviations under general graphs and costs. For the case of endogenous standard deviations, whether the nonconvexity of the social cost can be circumvented to obtain price of anarchy bounds for more general graphs and delay functions remains open.

7 Conclusions and Open Problems

We have set out to extend the classical theory of Wardrop equilibria and congestion games to the more realistic setting of uncertain delays, focusing on the methodology and questions of algorithmic game theory. The uncertainty of delays calls for models that incorporate players' attitudes towards risk. In this paper, we have focused on the model whereby players seek to minimize a linear combination of the expectation and standard deviation of delays along their chosen route.

The directions pursued in this work have opened a variety of questions which would be interesting to explore in future studies. Some of these questions are:

– What is the complexity of computing an equilibrium when it exists (exogenous standard deviations with atomic or nonatomic players; endogenous standard deviations with nonatomic players)?
– What is the complexity of computing the socially-optimal solution? What is the complexity of computing the socially-optimal flow decomposition if one knows the edge flow that represents a socially-optimal solution?
– Can there be multiple equilibria in the nonatomic game with endogenous standard deviations?

– What is the price of anarchy for stochastic Wardrop equilibria in the setting of nonatomic games with endogenous standard deviations, for general graphs and general classes of cost functions?
– Ordóñez and Stier-Moses considered the case of players with heterogenous attitudes toward risk [28]. Can some of the results in this paper be extended to that setting?

Of course, one could pursue other natural models and player utilities and build on or complement what we have developed here. In particular, our model might be enriched by also considering stochastic demands to make the demand side more realistic.

Acknowledgements We thank Hari Balakrishnan, José Correa and Costis Daskalakis for helpful discussions. The work of the first author was supported by the National Science Foundation grant number 0931550.

References

1. E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter. A survey on networking games in telecommunications. *Computers and Operations Research*, 33(2):286–311, 2006.
2. G. Andreatta and L. Romeo. Stochastic shortest paths with recourse. *Networks*, 18:193–204, 1988.
3. I. Ashlagi, D. Monderer, and M. Tennenholtz. Resource selection games with unknown number of players. In *Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 819–825, Hakodate, Japan, 2006. ACM Press, New York, NY.
4. M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT, 1956.
5. M. G. H. Bell and C. Cassir. Risk-averse user equilibrium traffic assignment: an application of game theory. *Transportation Research*, 36B(8):671–681, 2002.
6. D. P. Bertsekas and J. N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
7. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
8. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64:457–469, 2008.
9. J. R. Correa and N. E. Stier-Moses. Wardrop equilibria. In J. J. Cochran, editor, *Encyclopedia of Operations Research and Management Science*. Wiley, 2011.
10. S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the U.S. National Bureau of Standards*, 73B:91–118, 1969.
11. R. B. Dial. A probabilistic multi-path traffic assignment algorithm which obviates path enumeration. *Transportation Research*, 5(2):83–111, 1971.
12. Y. Y. Fan, R. E. Kalaba, and J. E. Moore. Arriving on time. *Journal of Optimization Theory and Applications*, 127(3):497–513, 2005.
13. Y. Y. Fan, R. E. Kalaba, and J. E. Moore. Shortest paths in stochastic networks with correlated link costs. *Computers and Mathematics with Applications*, 49:1549–1564, 2005.
14. H. Föllmer and A. Schied. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, 2004.
15. G. Langer. Traffic in the united states. *ABC News*, February 13 2005.

16. H. Liu, X. Ban, B. Ran, and P. Mirchandani. An analytical dynamic traffic assignment model with stochastic network and travelers' perceptions. *Transportation Research Record*, 1783:125–133, 2002.
17. H. K. Lo and Y.-K. Tung. Network with degradable links: capacity analysis and design. *Transportation Research*, 37B(4):345–363, 2003.
18. H. M. Markowitz. *Mean-Variance Analysis in Portfolio Choice and Capital Markets*. Basil Blackwell, Cambridge, MA, 1987.
19. P. B. Mirchandani and H. Soroush. Generalized traffic equilibrium with probabilistic travel times and perceptions. *Transportation Science*, 21:133–152, 1987.
20. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
21. J. F. Nash. Noncooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
22. Y. Nie and X. Wu. Shortest path problem considering on-time arrival probability. *Transportation Research Part B*, 43:597–613, 2009.
23. E. Nikolova. *Strategic algorithms*. PhD thesis, Massachusetts Institute of Technology. Dept. of Electrical Engineering and Computer Science, 2009.
24. E. Nikolova. Approximation algorithms for reliable stochastic combinatorial optimization. In *Proceedings of the 13th International Conference on Approximation*, APPROX/RANDOM'10, pages 338–351, Berlin, Heidelberg, 2010. Springer.
25. E. Nikolova, M. Brand, and D. R. Karger. Optimal route planning under uncertainty. In D. Long, S. F. Smith, D. Borrajo, and L. McCluskey, editors, *Proceedings of the International Conference on Automated Planning & Scheduling (ICAPS)*, pages 131–141, Cumbria, England, 2006.
26. E. Nikolova, J. A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *Lecture Notes in Computer Science 4168 (ESA 2006)*, pages 552–563, Springer, 2006.
27. N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK, 2007.
28. F. Ordóñez and N. E. Stier-Moses. Wardrop equilibria with risk-averse users. *Transportation Science*, 44(1):63–86, 2010.
29. C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, Hersonissos, Greece, 2001. ACM Press, New York, NY.
30. A. C. Pigou. *The Economics of Welfare*. Macmillan, London, 1920.
31. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
32. D. Schrank, T. Lomax, and S. Turner. Annual urban mobility report, 2010. Texas Transportation Institute. Available at `http://mobility.tamu.edu/ums`.
33. Y. Sheffi. *Urban Transportation Networks*. Prentice-Hall, Englewood, NJ, 1985.
34. Y. Sheffi and W. Powell. An algorithm for the traffic assignment problem with random link costs. *Networks*, 12:191–207, 1982.
35. M. Trahan. Probabilistic assignment: An algorithm. *Transportation Science*, 8(4):311–320, 1974.
36. S. V. Ukkusuri and S. T. Waller. Approximate analytical expressions for transportation network performance under demand uncertainty, 2009. Submitted.
37. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers, Part II, Vol. 1*, pages 325–378, 1952.