Columbia Business School

Decision, Risk & Operations
Working Papers Series

A Mean-Risk Model for the Stochastic Traffic Assignment Problem

E. Nikolova and N.E. Stier-Moses

May 2011; revised Dec. 2011
DRO-2011-03

A MEAN-RISK MODEL FOR THE STOCHASTIC TRAFFIC ASSIGNMENT PROBLEM

E. NIKOLOVA* AND N.E. STIER-MOSES**

* *Dept. of Computer Science and Engineering, Texas A&M University, College Station, TX, USA,* nikolova@tamu.edu

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

ABSTRACT. We embark on an agenda to investigate how stochastic travel times and risk aversion transform the traditional traffic assignment problem and its corresponding equilibrium concepts. Moving from deterministic to stochastic travel times with risk-averse users introduces non-convexities that make the problem more difficult to analyze. For example, even computing a best response of a user to the environment is still of unknown complexity [50, 48]. This paper focuses on equilibrium existence and characterization in the different settings of infinitesimal (non-atomic) vs. atomic users and fixed (exogenous) vs. congestion-dependent (endogenous) variability of travel times. We show that equilibria always exist in three out of the four possible combinations. The exception is the case of atomic users and endogenous variability. Because cost functions are non-additive (i.e., the cost along a path is *not* a sum of costs over edges of the path as it is assumed in the vast majority of network routing problems), solutions need to be represented as path flows since not all decompositions from edge to path-flows are equivalent. Nevertheless, we show that succinct representations of equilibria and optimal solutions always exist. Finally, we investigate the inefficiencies resulting from the stochastic nature of travel times. We obtain that under exogenous variability of travel times, the worst-case inefficiency of equilibria is exactly the same as when travel time functions are deterministic, meaning that in this case risk-aversion under stochastic travel times does not further degrade a system in the worst-case in relation to users' self-mindedness.

KEYWORDS: Non-additive Traffic Assignment Problem, Congestion Game, Stochastic Wardrop Equilibrium, Stochastic Nash Equilibrium, Risk aversion, Mean-standard deviation objective.

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 [64], "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." 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 that can be achieved when traffic is low. A common driver reaction in the face of heavy and uncertain traffic conditions is to look for alternate, sometimes longer but less crowded and more reliable routes [35]. With the widespread use of ever-improving technologies for measuring traffic, one might ask: what is a good routing strategy? And how does the risk-aversion of commuters transform the resulting traffic conditions?

We consider the traffic assignment problem on networks with stochastic travel times and analyze the resulting equilibria when strategic risk-averse commuters take into account the variability of their travel times. This approach generalizes the traditional model of Wardrop competition [71] by incorporating uncertainty.

Date: May 2011. This version: December 2011.

Risk aversion induces users to go beyond considering expected travel times. Since it is unlikely that they base their route-choice decisions on something as complicated as a full distribution of travel times along an exponential number of possible routes, it is reasonable that considering expected travel times and their standard deviations is a good first-order approximation on route selection. To incorporate the standard deviation of travel times into the users' objectives, we consider the traditional *mean-standard deviation* (mean-stdev) objective [26, 40] whereby users minimize the cost on a path, defined as its mean travel time plus a risk-aversion factor times the standard deviation of travel time along the same path. By linearity of expectations, the path mean equals the sum of edge means. However, the standard deviation along a path does not decompose as a sum over edges in the path because of the *risk-diversification effect*. Assuming independence of travel times on edges, the standard deviation equals the square root of the sum of the squared standard deviations on the edges of that path.

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 the travel time along it. Hence, we have a traffic assignment problem where users minimize a given percentile of their travel time, as opposed to the standard Wardrop model where they minimize the expected value of their travel time [53]. Note that this objective is also related to typical quantifications of risk, most notably the value-at-risk objective commonly used in finance, whereby one seeks to minimize travel time subject to arriving on time to a destination with at least, say, 95% chance.

There are alternative, simpler models that could be used to address uncertainty. The easiest would be to incorporate uncertainty on an edge-by-edge basis. This approach has been used in the past but it fails to exploit the risk-diversification effects of being exposed to multiple sources of uncertainty [69]. In our setting it would mean that standard deviations are *simply summed* along a path as opposed to taking the square root of the sum of squared standard deviations, as should be the case. Another alternative is to consider the mean-variance objective. This approach also reduces to a deterministic Wardrop equilibrium in which the travel time 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 and the objective seems less justified. In addition, this objective leads to solutions that are not intuitive in practice such as users selecting routes that are stochastically dominated by others. We highlight that this counterintuitive phenomenon may happen as well under the mean-stdev objective with some artificially constructed distributions, but it is guaranteed *not* to happen under normal distributions due to the equivalent percentile interpretation of the travel times along routes. In contrast, the mean-variance objective still suffers from this problem even in the case of normally-distributed uncertainty.

Readers familiar with the concept of coherent risk measures (see surveys [58, 34]) will recognize that the mean-standard deviation objective does not constitute a coherent risk measure, because it lacks monotonicity. Similarly, the value-at-risk measure (corresponding to the above-mentioned percentile objective) is not coherent because it is not convex. We remark that the assumptions underlying coherent risk measures were developed with respect to risk preferences in finance. While they may provide useful alternatives for the risk-averse objectives in the context of network games, they are not necessarily the only correct approach since routing preferences under uncertainty may differ axiomatically from preferences in portfolio optimization. For example, an intercity bus that needs to observe a given departure and arrival schedule may prefer a more certain (albeit dominated) path if there are constraints that prevent early arrival such as the lack of parking space at the bus depot. Thus, while monotonicity is a reasonable requirement for risk measures in the context of finance, it may sometimes be dropped in networks applications such as transportation and telecommunications.

We assume that the expected travel time and the standard deviation of travel time are nondecreasing functions that depend on the load of an edge.[1] To provide an example, one gets increasing functions on the load when each segment of a network represents a queue. In that case, the expected travel time and standard deviation are increasing functions on the load of the queue. For tractability reasons, we also consider a simpler case in which the standard deviation is considered exogenous to the model and hence independent of the flow. Although this may seem simplistic, it is a common simplification done in earlier work in stochastic network models [12, 53]. Furthermore, the traditional regression models commonly used in a multitude of disciplines make a similar simplification when they assume that the error term has a fixed standard deviation independent of the explanatory variables. As we will see below, without this simplification the standard deviation is endogenously determined, which can cause equilibria to fail to exist.

Even in the simpler case of exogenous standard deviations, the form of path costs is complicated by the square root used to compute the standard deviation of the paths. In this setting, solving a user'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. To the best of our knowledge, a precise characterization of the complexity of the subproblem is open; the best algorithms known so far run in time $n^{O(\log n)}$ for networks of n nodes, which is between polynomial and exponential [50, 48]. This is in sharp contrast to the subproblem of the Wardrop network game—a shortest path problem, which admits efficient solutions such as Dijkstra's algorithm [22].

Our mean-stdev model works for *arbitrary* distributions with finite first and second moment. To simplify the analysis, throughout this paper we assume that travel times in different edges are uncorrelated. Nevertheless, correlation is to be expected in practice; e.g., 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 [47]. It is possible to obtain a graph with independent travel times on edges where all our results and algorithms carry through. Nie and Wu also consider local correlations [45].

Summary of Results. We generalize the traditional model of Wardrop competition [71] by incorporating stochastic travel timess (see Section 3 for the model, and Section 2 for related research). Technically, this model is much harder to analyze than the traditional one because it is *non-additive*, namely the cost of a path is not equal to the sum of costs of edges along the path [28]. This in turn means that an equilibrium in the stochastic setting does not decompose to equilibria in subnetworks of the given network, leading to computational and structural complications. Depending on the specific details of the application one has in mind, users may be small or large [29]. We consider both infinitesimal users, referred to as the *non-atomic* case, as well as users that control a strictly positive demand, referred to as the *atomic* case.

To analyze the problem and to establish the existence of equilibrium, we draw from a diverse spectrum of tools from potential games and convex analysis to the theory of variational inequalities and nonconvex

[1] Actually, some of our results also extend to the non-separable case, where these functions depend on the full vector of loads of all edges of the network, but this will not be the focus of this study.

	Exogenous Standard Deviations	Endogenous Standard Deviations
Nonatomic Users	Equilibrium exists (It solves exponentially-large convex program)	Equilibrium exists (It solves variational inequality)
Atomic Users	Equilibrium exists (Game is potential)	No pure strategy equilibrium

TABLE 1. Existence of equilibria in mean-risk stochastic traffic assignment problems.

(stochastic) shortest paths. We consider four settings of nonatomic vs. atomic users and exogenous (Section 4) vs. endogenous variability of travel times (Section 5). Our conclusions and methods are different in each of these settings. In the nonatomic case with standard deviation of travel times given exogenously, we prove that equilibria always exist using a convex problem with exponentially-many variables similar to that of Ordóñez and Stier-Moses [53]. The atomic case with exogenous standard deviations is shown to be a potential game and therefore a pure-strategy Nash equilibrium always exists. To characterize the equilibria of the nonatomic version of the problem when the standard deviations of travel times are endogenous, we use a variational inequality formulation [31, 67, 18] that draws ideas from the nonlinear complementary problem formulation of Aashtiani and Magnanti [1]. In this case, an equilibrium always exists; in fact, not only for our specific mean-stdev objective but also for any general continuous objective. In contrast, the atomic case with endogenous standard deviation does not always admit a pure-strategy Nash equilibrium. We summarize these results in Table 1.

Next, we investigate if there is a succinct representation (in terms of a small set of paths) of user and system optimal flows in the case of non-atomic users with stochastic travel times (Section 6). Our results here are independent of whether the standard deviations are exogenous or endogenous. We prove that if one is given a solution (either a Wardrop equilibrium or a system optimum) as an edge-flow, not every path decomposition is a solution, in contrast to the deterministic case where every decomposition works. Nevertheless, there is always a succinct solution that uses at most $|E| + |K|$ paths, where E is the set of edges in the network and K is the set of origin-destination pairs. Although the complexity of computing a solution is left open (actually, even the complexity of computing a single stochastic shortest path is open), this result says that there is some hope because at least solutions can be efficiently encoded.

Finally, we quantify the inefficiency of mean-risk Wardrop equilibria under stochastic travel times with respect to the socially-optimal solution, for the case of nonatomic users (Section 7). The social optimum is defined as the flow minimizing the total cost incurred by users, as given by their mean-stdev objective. Surprisingly, under exogenous standard deviations, uncertainty and risk aversion do not exacerbate the inefficiency of equilibria. The price of anarchy remains equal to that of deterministic nonatomic games. Namely, it is $4/3$ for the case of linear expected travel times [61] and $(1 - \beta(\mathcal{L}))^{-1}$ for an appropriately defined constant $\beta(\mathcal{L})$ for expected travel time functions in a class $\mathcal{L}$ [60, 15, 16].

The case of endogenous standard deviations presents a significant additional difficulty that makes the square root terms in different paths interrelated functions of the path flow that cannot be analyzed separately; a general price of anarchy bound for this case remains elusive. Nevertheless, we show that, despite the square root term, the path costs are convex whenever the individual travel times and standard deviations on edges are convex. Consequently, we present sufficient conditions for convexity of the social cost, which are similar to the sufficient conditions for uniqueness of equilibrium in its variational inequality characterization. Unfortunately, these conditions are fragile and in general the social cost will not be convex and may admit a non-connected set of multiple global minima but we can still identify settings where the price of anarchy is 1.

We provide concluding remarks and some open questions in Section 8.

2. 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 users jointly select shortest routes [71]. The model was formalized in an influential book by Beckmann *et al.* where they lay out the mathematical foundations to analyze network games [7]. These models find applications in various application domains such as in transportation [65] and telecommunication [3] networks. In the last decade, these types of models have received renewed attention with many studies aimed at understanding under what conditions these games admit an equilibrium, what uniqueness properties are satisfied by these equilibria, what methods can be used to compute equilibria efficiently, what price is paid for having competition instead of a centralized solution, and what are good ways to align incentives so an equilibrium becomes socially optimal. For references on these topics from a perspective similar to ours, we refer the readers to the surveys [17, 52].

In the majority of models used by theoreticians who study the properties of equilibria in networks, and by practitioners who compute solutions to real problems, travel times have been considered deterministic. For instance, most of the previous work assumes that travel times depend on the load of the edges, with different degrees of generality. In recent times, researchers progressively started paying more attention to risk aversion and began incorporating various forms of uncertainty to their models (see, e.g., [5, 37, 41, 70] and some more references below). Nevertheless, 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 [21], which has been studied in detail and used in practice extensively [66, 68]. Under it, different users *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." Daganzo and Sheffi [20] looked at the case of dependent route costs, while Fisk [25] studied the model in the context of congested networks, obtaining an equivalent optimization problem. Methods that avoid route enumeration have been proposed by Bell [8], Larsson, Liu, and Patriksson [36], and Maher [39], also leading to equivalent optimization problems in the spirit of Fisk's. Based on Akamatsu [2], Baillon and Cominetti [6] proposed a more general concept called *Markovian traffic equilibrium*, provided an equivalent optimization problem and established the convergence of the method of successive averages in that context. But the bottom line is that this model considers that *perceptions* on different routes are stochastic, and not that the travel times themselves are. For this reason, the model presented in this work is complementary to the stochastic user equilibrium approach.

The route-choice model in this paper consists of users that select the path that minimizes the mean plus a multiple of the standard deviation of travel time. This problem belongs to the class of stochastic shortest path problems (we refer the reader to some classic references [4, 11] and some newer ones [23, 24, 49, 46]). Wu and Nie [72] make use of stochastic dominance to characterize admissible paths in a route choice model with uncertain travel times. Besides stochastic formulations, there have been other approaches to this problem. For example, Bell and Cassir consider that travel times are set by an adversary who will pick the worst-possible travel time for the user [9], and Bertsimas and Sim propose a robust optimization approach that considers a budget of uncertainty that limits the number of edges on which actual travel times are different from the mean travel times [12].

Going from route choice into equilibrium network assignment problems, Lo and Tung study a probabilistic user equilibrium model for networks with stochastic capacity [38]. Their 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 model most related to our work is that of Ordóñez and Stier-Moses [53]. 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 travel times along flow-bearing paths are minimal. They also use their algorithm to compare equilibria with risk-averse players to those with risk-neutral players, as in the standard Wardrop model. 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 the atomic case where users control a positive amount of flow. Following up on Ordóñez and Stier-Moses, Nie also studies percentile equilibria [44]. He studies an instance with two edges and exogenous standard deviations in detail, provides a gradient projection algorithm to find percentile equilibria, and uses it to perform a computational study. Like us, he also considers congestion-dependent standard deviations of travel times.

3. The Model

We consider a directed graph $G = (V, E)$ with an aggregate demand of d_k units of flow between origin-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. The users in the network—i.e., the players of the game—must choose routes that connect their origins to their destinations. We encode the collective decisions of users in 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 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 network is subject to congestion, modeled with stochastic travel time functions $\ell_e(x_e) + \xi_e(x_e)$ for each edge $e \in E$. Here, $\ell_e(x_e)$ measures the expected travel time when the edge has flow x_e, and $\xi_e(x_e)$ is a random variable that represents a noise term on the travel time, encoding the error that $\ell_e(\cdot)$ makes. The function $\ell_e(\cdot)$ is assumed continuous and non-decreasing. The random variable $\xi_e(x_e)$ has expectation equal to zero and standard deviation equal to $\sigma_e(x_e)$, for a continuous and non-decreasing function $\sigma_e(\cdot)$. Although the distribution generally depends on the flow value x_e, we will separately consider the simplified case in which the function $\sigma_e(x_e)$ is a constant σ_e given exogenously, and therefore independent from x_e. We also assume that these random variables are all uncorrelated with each other. As explained in the introduction, risk-averse players choose paths according to the mean-standard deviation (mean-stdev) objective, a linear combination of the expectation and standard deviation of the travel time along the route. For simplicity, throughout the paper we refer to the mean-stdev objective as the cost along a route. Formally, the cost along route π is

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

where $\gamma \geq 0$ is a constant that quantifies the user risk-aversion, which we assume homogeneous.

The *nonatomic* version of the problem considers the setting where there are an infinite number of users who control an insignificant amount of flow each so that the path choice of a single user does not unilaterally affect costs experienced by others (even though the joint actions of several players affect other players). The following definition captures that users at equilibrium 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* $k \in K$ *and for every path* $\pi \in \mathcal{P}_k$ *with positive flow,* $Q_\pi(\mathbf{f}) \leq Q_{\pi'}(\mathbf{f})$ *for any* $\pi' \in \mathcal{P}_k$ *.*

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. When users control a positive demand (the atomic case), there are two possibilities: in the splittable case users 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 re-routing the flow to another path.

Definition 2. *A pure-strategy* stochastic Nash equilibrium *of the atomic unsplittable routing game is a flow* $\mathbf{f}$ *such that for every* $k \in K$ *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 any* $\pi' \in \mathcal{P}_k$ *. Here,* $\mathcal{I}_\pi$ *denotes a vector that contains a one for path* π *and zeros otherwise.*

This game always admits a mixed strategy equilibrium because it is a finite, normal form game [43]. Nevertheless, we focus on the existence of a pure-strategy Nash equilibrium because it is a more natural solution concept.

To quantify the quality of solutions, and in particular of equilibria, we define a social cost function that will allow us to compare different flows and determine the optimal one. We adopt a natural social-cost function, given by the total cost among all users:

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

4. Exogenous Standard Deviations

In this section, we consider that noise factors affecting travel times are exogenous, which result in constant standard deviations $\sigma_e(x_e) = \sigma_e$ that do not depend on the flow on the edge. The motivation for studying this case is given by variations of travel time that depend on external factors such as the weather, events, traffic signals, or other phenomena that change the road capacity independently of the flow (for more details, see Ordóñez and Stier-Moses [53]). Although this setting is more restrictive, it constitutes a first step in understanding how variability of travel times influences equilibrium models. In the next section, we study the more general setting where the standard deviation function may depend on the flow.

For constant standard deviations of travel times, the path cost (1) can be written as $Q_\pi(\mathbf{f}) = \sum_{e\in\pi} \ell_e(f_e) + \gamma(\sum_{e\in\pi} \sigma_e^2)^{1/2}$. It is important to highlight that the second term is a constant that depends on the path but does not depend on the flow on the edges. We investigate the existence of equilibria and provide a characterization, first for the nonatomic case and then for the atomic one.

4.1. **The Nonatomic Case.** Despite the challenge posed by the non-additive cost function $Q_\pi(\cdot)$, we show that an equilibrium always exists using a path-based convex programming formulation given by Ordóñez and Stier-Moses [53].

Theorem 4.1. *A nonatomic routing game with exogenous standard deviations always has a stochastic Wardrop equilibrium.*

Proof. Even though we cannot separate the cost into a sum of costs over the edges as traditional formulations of Wardrop equilibria [7], we can characterize the equilibrium using a convex program as follows:

$$
\begin{aligned}
\min & \sum_{e\in E} \int_0^{x_e} \ell_e(z)dz + \sum_{\pi\in\mathcal{P}} \gamma f_\pi \sqrt{\sum_{e\in\pi} \sigma_e^2} \qquad (3)\\
\text{s.t. } & x_e = \sum_{\pi\in\mathcal{P}:\, e\in\pi} f_\pi \quad \text{for } e\in E,\\
& d_k = \sum_{\pi\in\mathcal{P}_k} f_\pi \quad \text{for } k\in K,\\
& f_\pi \geq 0 \quad \text{for } \pi\in\mathcal{P}\,.
\end{aligned}
$$

The term in the objective with the square root is linear in the flow implying that the objective is a continuous convex function provided the functions $\ell_e(x)$ are continuous and nondecreasing. Since the constraint set of feasible flows is a polytope, a minimum is attained. Moreover, the first order conditions for the convex program exactly match the definition of Wardrop equilibrium, proving its existence. □

The formulation in the previous proof also implies that the equilibrium is unique when the convex objective function (3) is strictly convex, which leads to the following corollary.

Corollary 4.2. *The stochastic Wardrop equilibrium of the nonatomic routing game with exogenous standard deviations is unique (in terms of edge loads) whenever the expected travel time functions $\ell_e(\cdot)$ are strictly increasing.*

Besides proving existence, the formulation (3) also provides a way to compute this equilibrium using a column generation procedure. The convex program (3) contains exponentially many variables (the flows on all paths) and a polynomial number of constraints. We will see in Section 6 that an equilibrium always has a succinct decomposition that uses at most $|E| + |K|$ paths; unfortunately, since we do not know ahead of time which paths these are, we cannot write a succinct version of the convex program. Nevertheless, this succinctness property provides a practical method for computing equilibria. We refer the reader to Ordóñez and Stier-Moses [53] for details on computation and for a computational study of the equilibria of these games.

In the case of constant expected travel times, the objective (3) coincides with the social cost, and both problems reduce to computing a stochastic shortest path for each origin destination pair. Thus, both the equilibrium and social optimum computation are at least as hard as the stochastic shortest path problem [50, 48].

Theorem 4.3. *When the expected travel times and standard deviations are constant for each edge, the equilibrium and social optimum coincide and can be found in time* $n^{O(\log n)}$.

Proof. In the case of constant expected travel times, say $\ell_e(x_e) = b_e \geq 0$ for all edges e, the convex program formulation (3) is the same as the social optimum minimization program and turns into the following linear program:

$$\begin{aligned} &\min \sum_{\pi \in \mathcal{P}} f_\pi \left(\sum_{e \in \pi} b_e + \gamma \sqrt{\sum_{\pi \in \mathcal{P}} \sigma_e^2} \right) \\ &\text{s.t. } d_k = \sum_{\pi \in \mathcal{P}_k} f_\pi \quad \text{for } k \in K, \\ &\quad f_\pi \geq 0 \quad \text{for } \pi \in \mathcal{P}\,. \end{aligned} \tag{4}$$

The optimal solution to this problem assigns all demand between a origin-destination pair on the least-cost path between that pair. So, for each origin-destination pair $k \in K$, we solve the stochastic shortest path problem $\min_{\pi \in \mathcal{P}_k} \left\{ \sum_{e \in \pi} b_e + \gamma (\sum_{e \in \pi} \sigma_e^2)^{1/2} \right\}$. Since $|K| \in O(n^2)$, the problem can be solved in time $n^{O(\log n)}$. □

4.2. **The Atomic Case.** Now, we switch our attention to atomic unsplittable routing games and show that they admit a potential function. We prove this using the characterization given by Monderer and Shapley [42]. This property implies that a pure-strategy stochastic Nash equilibrium always exists.

Theorem 4.4. *An atomic unsplittable routing game with exogenous standard deviations always has a pure-strategy stochastic Nash equilibrium.*

Proof. Since the game is unsplittable and players control a unit demand each, a flow is described by a set of paths $\pi := (\pi_i)_{i \in K}$ chosen by players. The corresponding edge-flow is denoted by f^π; that is, f_e^π counts how many players selected a route that includes edge e. Finally, the set $-J$ refers to the complement of the players in a given set J. Following the characterization of potential games by Monderer and Shapley [42], we consider the strategy graph that associates a node to every vector of players' strategies. This graph contains an edge between two nodes whenever their corresponding vectors of strategies differ exactly in the strategy of a single player. To prove that a potential function exists, it suffices to show that the total change of players' costs is zero along an arbitrary cycle of length four in the strategy graph.

Let us consider two players i and j, who initially select routes π_i and π_j, respectively. The cycle of length four must consist of the following four moves in the strategy graph, where both players select a new route π'_i and π'_j, respectively. Indeed, the cycle consists of vertices π, $\pi_{(2)} = (\pi'_i, \pi_j, \pi_{-\{i,j\}})$, $\pi_{(3)} = (\pi'_i, \pi'_j, \pi_{-\{i,j\}})$, $\pi_{(4)} = (\pi_i, \pi'_j, \pi_{-\{i,j\}})$, and back to π. Let us now evaluate the cost variations. When player i changes its strategy from π_i to π'_i, his cost difference is

$$Q_i(\pi_{(2)}) - Q_i(\pi) = \sum_{e \in \pi'_i} \ell_e(f_e^{\pi_{(2)}}) + \gamma \sqrt{\sum_{e \in \pi'_i} \sigma_e^2} - \sum_{e \in \pi_i} \ell_e(f_e^\pi) - \gamma \sqrt{\sum_{e \in \pi_i} \sigma_e^2}\,.$$

When player j changes its strategy from π_j to π'_j, his cost difference is

$$Q_j(\pi_{(3)}) - Q_j(\pi_{(2)}) = \sum_{e\in\pi'_j} \ell_e(f_e^{\pi(3)}) + \gamma\sqrt{\sum_{e\in\pi'_j}\sigma_e^2} - \sum_{e\in\pi_j}\ell_e(f_e^{\pi(2)}) - \gamma\sqrt{\sum_{e\in\pi_j}\sigma_e^2}\,.$$

When player i changes its strategy back from π'_i to π_i, his cost difference is

$$Q_i(\pi_{(4)}) - Q_i(\pi_{(3)}) = \sum_{e\in\pi_i} \ell_e(f_e^{\pi(4)}) + \gamma\sqrt{\sum_{e\in\pi_i}\sigma_e^2} - \sum_{e\in\pi'_i}\ell_e(f_e^{\pi(3)}) - \gamma\sqrt{\sum_{e\in\pi'_i}\sigma_e^2}\,.$$

Finally, when player j changes its strategy back from π'_j to π_j, his cost difference is

$$Q_j(\pi) - Q_j(\pi_{(4)}) = \sum_{e\in\pi_j} \ell_e(f_e^{\pi}) + \gamma\sqrt{\sum_{e\in\pi_j}\sigma_e^2} - \sum_{e\in\pi'_j}\ell_e(f_e^{\pi(4)}) - \gamma\sqrt{\sum_{e\in\pi'_j}\sigma_e^2}\,.$$

Summing the previous equations, all the terms with square roots cancel out, leading to

$$\widehat{Q_i}(\pi_{(2)}) - Q_i(\pi) + \widehat{Q_j}(\pi_{(3)}) - Q_j(\pi_{(2)}) + \widehat{Q_i}(\pi_{(4)}) - Q_i(\pi_{(3)}) + \widehat{Q_j}(\pi) - Q_j(\pi_{(4)})\,,$$

where $\widehat{Q}$ denotes a modified path cost that ignores the variability of travel times. This implies that the change in the original cost summed across players is equal to zero because the associated deterministic routing game with costs $\ell_e(x)$ is potential [59]. □

It is well known that equilibria for deterministic, nonatomic games are essentially unique (Corollary 4.2 is a generalization to that). This means that if there is more than one equilibrium, the travel time along any edge is the same under different equilibria. Consequently, users experience the same cost under all equilibria, and different equilibria are indistinguishable from each other, both from the users' perspective and from the edges' perspective. Instead, there can be different pure-strategy Nash equilibria in the deterministic, atomic game. A simple example that illustrates that is given by three parallel edges with (deterministic) travel time functions equal to $x + 1$, $x + 1$ and $x + 1/2$, respectively, and two players that control a unit demand each. When none of the players select edge 3, one of them can profitably deviate to it, so a necessary condition for equilibrium is that exactly one player selects edge 3. Hence, there are four equilibria in total, where one player selects edge 1 or 2 and the other player selects edge 3. At equilibrium players can experience a travel time of 2 or $3/2$ depending on the selected edge. Furthermore, edge 1 may have a travel time of 1 or 2 at equilibrium, depending whether a player selected it or not. The multiplicity of equilibria arises because of the atomic nature of the players, and not because of the uncertain travel times.

5. Endogenous Standard Deviations

In this section, we consider the more general case of flow-dependent standard deviations of travel times. This makes the standard deviations endogenous to the game, and can be used to model the impact of congestion not only on the expectation of travel time but also on the ensuing variability. For instance, incidents are more likely when there is more traffic in a road. Other sources of delay that are also more likely to occur are that more users may be looking for parking, may double-park and may stop to pickup or drop off passengers. Using a variational-inequality formulation in the space of path flows, we show that in the nonatomic case equilibria continue to exist. Unfortunately, contrary to the setting with deterministic travel times, a formulation using a minimization problem is not possible. In the case of atomic users, equilibria may fail to exist as a consequence of the consideration of endogenous standard deviations of travel times.

We start with an example that illustrates how an equilibrium changes when standard deviations are endogenous. Assume a demand of $d = 1$ and consider a network consisting of two parallel edges with travel times $\ell_1(x) = x$ and $\ell_2(x) = 1$, followed by a chain of k edges that users must traverse. This instance admits two paths, each comprising one of the two parallel edges and the chain. We let L denote the expected travel time 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 the classic instance with two routes put forward by Pigou [56, 61], the equilibrium with endogenous standard deviations changes significantly. Indeed, it can be characterized by the roots of the degree-4 polynomial $(1 - 4s^2)x^4 + 4s^2x^3 + (4s^4 - 2s^2 - 4ks^2)x^2 - 4s^4x + s^4$ that are in $[0, 1]$, where x denotes the flow on one of the two paths. Although in principle it is not evident whether these roots exist or not, we will see that an equilibrium always does.

An important insight that arises 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 subgraphs: this is a major challenge to designing efficient algorithms for computing equilibria, or even just best responses, for which traditional dynamic-programming type approaches will likely fail (e.g., Dijkstra's algorithm [22]).

5.1. **The Nonatomic Case.** At the end of the 1970's, Smith [67] and Dafermos [18] proposed to characterize Wardrop equilibria in nonatomic routing games using the solutions of variational inequality problems. Some earlier research used this approach for nonadditive models like ours [28, 1, 53]. All these papers show that a flow minimizes a modified cost function (2) that *holds path costs fixed* if and only if the flow is at equilibrium. In other words, $\mathbf{f}$ is at equilibrium if for any feasible flow $\mathbf{f}'$,

$$\mathbf{Q}(\mathbf{f}) \cdot (\mathbf{f} - \mathbf{f}') \leq 0, \tag{5}$$

where $\mathbf{Q}(\mathbf{f})$ denotes the vector of costs along all paths $(Q_\pi(\mathbf{f}))_{\pi \in \mathcal{P}}$. Hence, the existence of an equilibrium can be proved through establishing the existence of a solution to the variational inequality above. This is a corollary of the following result from the theory of variational inequalities because the set of feasible flows is convex and compact, and the cost function is continuous. Note that the existence of an equilibrium also holds in the much more general setting where the travel time functions depend not only on the flow of the given edge but also on the other edges, as long as this dependence is continuous. This is referred to as *nonseparable* travel time functions in the literature.

Theorem 5.1. [31] *Let $\mathbb{K} \subset \mathbb{R}^N$ be a compact convex set and let $\mathbf{Q} : \mathbb{K} \to \mathbb{R}^N$ be a continuous mapping. Then, there exists a vector $\mathbf{x} \in \mathbb{K}$ such that $\mathbf{Q}(\mathbf{x}) \cdot (\mathbf{x} - \mathbf{y}) \leq 0$ for all $\mathbf{y} \in \mathbb{K}$.*

Corollary 5.2. *The nonatomic routing game with endogenous standard deviations has a stochastic Wardrop equilibrium.*

In contrast to the case of exogenous standard deviations, however, the game with endogenous standard deviations is not *potential* [42] and equilibria cannot be easily characterized as the solution to a (convex) optimization problem.

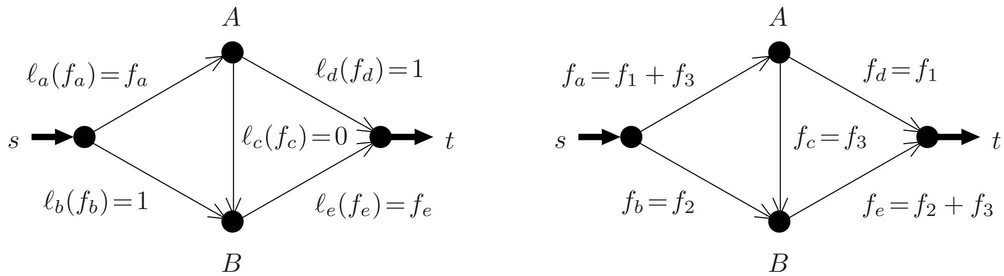

(a) Expected travel time functions.

(b) Total flows in each edge.

FIGURE 1. An example that shows that there is no cardinal potential function.

Proposition 5.3. *The stochastic routing game with endogenous standard deviations does not admit a cardinal potential.*

Proof. For games with infinite player sets, as is the nonatomic routing game, Sandholm provides a condition that characterizes potential games. This condition, called *externality symmetry* (which we define below), does not hold in our game. To see this, suppose that our game admits a cardinal potential function $\Phi : \mathbb{R}^{|\mathcal{P}|} \to \mathbb{R}$. Note that the domain is the set of all path flows, not only those that satisfy demands (technically, this is called a *full population game*; see, e.g., [63]). By definition, this is a continuously differentiable function whose gradient is the vector of path-cost functions. Equivalently, since the path-cost functions are smooth, they must satisfy the *externality symmetry* condition [62], which means that for any two paths $\pi, \pi' \in \mathcal{P}$, the effect on the cost of path π' of adding flow on path π is equal to the effect on the cost of path π of adding flow on path π'. In other words, the cross partial derivatives are the same when the order is exchanged:

$$\frac{\partial^2 \Phi(\mathbf{f})}{\partial f_\pi \partial f_{\pi'}} = \frac{\partial Q_\pi(\mathbf{f})}{\partial f_{\pi'}} = \frac{\partial Q_{\pi'}(\mathbf{f})}{\partial f_\pi} = \frac{\partial^2 \Phi(\mathbf{f})}{\partial f_{\pi'} \partial f_\pi}\,.$$

However, the following example based on the Braess paradox network shows that externality symmetry is not satisfied, proving the claim.

Consider the travel time functions indicated in Figure 1 with standard deviation functions equal to f_e for all edges e. There are three possible paths: top, down and zigzag, with flows denoted by f_1, f_2, f_3 and cost functions respectively equal to

$$\begin{aligned}
Q_1(\mathbf{f}) &= 1 + f_1 + f_3 + \sqrt{(f_1+f_3)^2 + f_1^2},\\
Q_2(\mathbf{f}) &= 1 + f_2 + f_3 + \sqrt{(f_2+f_3)^2 + f_2^2},\\
Q_3(\mathbf{f}) &= f_1 + f_2 + 2f_3 + \sqrt{(f_1+f_3)^2 + f_3^2 + (f_2+f_3)^2}.
\end{aligned}$$

Considering the cross effects of paths 1 and 3:

$$\frac{\partial Q_1(\mathbf{f})}{\partial f_3} = 1 + \frac{f_1+f_3}{\sqrt{(f_1+f_3)^2+f_1^2}} \qquad \text{and} \qquad \frac{\partial Q_3(\mathbf{f})}{\partial f_1} = 1 + \frac{f_1+f_3}{\sqrt{(f_1+f_3)^2+f_3^2+(f_2+f_3)^2}}\,.$$

we see that $\partial Q_1(\mathbf{f})/\partial f_3 \neq \partial Q_3(\mathbf{f})/\partial f_1$, so externality symmetry does not hold. □

5.2. **Uniqueness.** As in the deterministic case, the stochastic routing game may have multiple flows that are at equilibrium when expected travel times and their standard deviations are not strictly increasing with flow. (We present one example below in Lemma 6.1, which shows that under constant expected travel times and standard deviations, there are different edge flows at equilibrium. In the example, any flow of the form $(f, 1-f, 1-f, f)$ for $f \in [0,1]$ is an equilibrium.) Thus, a relevant question is whether there exists a unique equilibrium when the expected travel time and/or standard-deviation functions are strictly increasing, as is the case when travel times are deterministic. Although a unique equilibrium exists for extreme risk attitudes, we leave the question for general risk attitudes open. We show how some of the standard methods that are used for characterizing equilibria and establishing uniqueness using the theory of variational inequalities and nonlinear complementarity problems fail. The following is a classical result usually employed to settle questions of this kind (see, e.g., [30]).

Theorem 5.4. *Consider the variational inequality $F(x) \cdot (x-y) \leq 0$ for all $y \in X$, over a nonempty, compact and convex domain $X \subset \mathbb{R}^n$. If the mapping $F : \mathbb{R}^n \to \mathbb{R}^n$ is* strictly monotone *over X, meaning that*

$$[F(x) - F(y)] \cdot (x-y) > 0 \qquad \forall x \neq y \in X\,, \tag{6}$$

then the variational inequality has at most one solution.

There are some other weaker notions, such as monotonicity and pseudo-monotonicity, that can be used to prove existence of solutions and other properties but they are not enough to guarantee uniqueness.

Remark 5.5. *It is easy to see that the sum (and also the convex combination) of two monotone operators is monotone. In our setting, the path-cost operator is a linear combination of the mean and the standard deviation of travel times along paths. The mean is always monotone because it is separable and non-decreasing. If the standard deviations are monotone on a class of graphs, we would directly obtain that path costs are monotone, resulting in uniqueness of equilibrium (under appropriate conditions for* strict *monotonicity such as strictly increasing path means). Conversely, for monotonicity to fail, it needs to fail in the case of* infinitely risk-averse users*, whereby path costs are given only by the standard-deviation term.*

Following the remark, we present counterexamples for monotonicity under the case of infinitely risk-averse users where $\gamma \to \infty$.[2] The key insight is that the square-root function is not monotone. By continuity, these counterexamples can be extended to the case with moderately risk-averse users where costs include positive expectation terms. We say that a mapping is *pseudo-monotone* over a domain X if

$$F(y) \cdot (x-y) \geq 0 \text{ implies } F(x) \cdot (x-y) \geq 0 \qquad \forall x, y \in X. \tag{7}$$

We present the following counterexample for the most general definition of monotonicity because it implies that none of the other monotonicity properties hold for stochastic routing games with endogenous standard deviations. In particular, the operator is not monotone, nor strictly monotone.

Proposition 5.6. *The path-cost operator of the nonatomic routing game with endogenous standard deviations is not pseudo-monotone.*

[2] An easy way to capture an infinite risk-aversion is by assuming that mean travel times are zero. Note that the possibility of having negative realizations of travel times does not impose limitations since one could add an appropriate constant to all edges without changing the solutions of the examples that we provide.

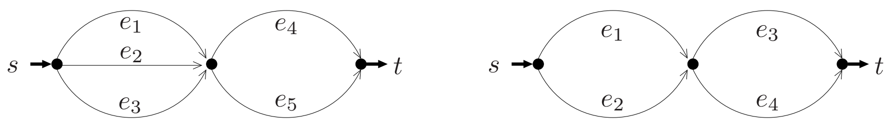

(a) Pseudo-monotonicity (b) Monotonicity

FIGURE 2. Counterexamples to (pseudo-)monotonicity.

Proof. Consider the network on the left of Figure 2, and assume that the expectation of travel time in each edge is zero and that the standard deviation $\sigma_e(f_e)$ on the edge equals f_e. Hence, $Q_\pi(\mathbf{f}) = (\sum_{e\in\pi} f_e^2)^{1/2}$. We refer to path (e_i, e_j) in this network by π_{ij} and we encode the full vector of flows as $\mathbf{f} = (f_{14},\ f_{24},\ f_{34},\ f_{15},\ f_{25},\ f_{35})$. Flow vectors $\mathbf{f} = (0, 0, 0.1, 0.2, 0.7, 0)$ and $\mathbf{f}' = (0.1, 0, 0, 0, 0.7, 0.2)$ violate pseudo-monotonicity because $Q(\mathbf{f}') \cdot (\mathbf{f} - \mathbf{f}') = 0.00494$ and $Q(\mathbf{f}) \cdot (\mathbf{f} - \mathbf{f}') = -0.00494$. □

An even simpler instance can be used to provide a counterexample to monotonicity. Indeed, consider the network on the right of Figure 2 with similar characteristics as that in the previous proposition. Denoting flows as $\mathbf{f} = (f_{13}, f_{23}, f_{14}, f_{24})$, flow vectors $\mathbf{f} = (0, 0.1, 0.2, 0.7)$ and $\mathbf{f}' = (0.1, 0, 0, 0.9)$, violate monotonicity because $(\mathbf{f} - \mathbf{f}') \cdot [Q(\mathbf{f}) - Q(\mathbf{f}')] = -0.00114$.

Although Proposition 5.6 implies that strict monotonicity does not hold, we can nonetheless prove uniqueness when users have extreme risk attitudes. In particular, we show that in those cases the stochastic game resembles a deterministic one.

Proposition 5.7. *The equilibrium of a stochastic routing game with endogenous standard deviations is unique in the two extreme settings where users are either risk-neutral or infinitely risk-averse, for strictly increasing expected travel times and standard deviations.*

Proof. When users are risk-neutral, the stochastic game trivially reduces to the deterministic game with travel times given by the expected travel time functions. Thus, we already know that the equilibrium is unique [7]. When users are infinitely risk-averse, although we saw that the standard deviation operator is not monotone, the game in which path costs are squared admits the same equilibria as the original game. This transformation makes the path costs additive and hence the game has a unique equilibrium because it is equivalent to a deterministic one with increasing travel times (given by the variance functions on each edge). □

5.3. **The Atomic Case.** In contrast to the nonatomic case, the atomic unsplittable routing game may not have pure-strategy Nash equilibria. Note that a mixed-strategy Nash equilibrium always exists (under the standard expected cost for mixing) since there are a finite number of players and strategies [43].

Proposition 5.8. *The atomic unsplittable routing game with endogenous standard deviations may not have pure-strategy Nash equilibria, even in the case of a single source, a single sink and a series parallel network with affine cost functions.*

Proof. Consider two users that want to route one unit of flow from s to t in the graph shown in Figure 3. Let the mean travel time and standard deviation on edges e_1, e, e_2, e_3 be $(4.8, 0)$, $(x, \sqrt{2}x)$, $(x, 1)$, $(x + 0.4, 0)$, respectively. We refer to the three paths in the graph as π_1, π_2, π_3, where path π_i uses edge e_i. The table in

Figure 3 shows the path costs under each strategy for risk-aversion coefficient $\gamma = 1$. A simple inspection shows that the game does not have a pure-strategy Nash equilibrium. $\square$

6. Succinct Representations of Solution

We now turn our attention to how one can decompose equilibria and socially optimal solutions represented as edge-flow vectors into path-flow vectors. Furthermore, the hope of efficient algorithms to compute those solutions depends on the existence of succinct vectors of path-flows, meaning that not too many paths are used. In this section, we set to study these questions, exclusively from the perspective of the nonatomic routing game.

Decompositions are easy in deterministic routing games: any path-flow decomposition of an equilibrium or a social optimum, given as an edge-flow, 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 in different path costs. In particular, for an equilibrium or a system optimum, given edge-flows, some path-flow decompositions are at equilibrium or optimal, respectively, and others are not.

The next lemma illustrates that shortest paths with respect to our nonadditive path costs do not need to satisfy Bellman equations since a subpath of a shortest path need not be shortest.

Lemma 6.1. *In a nonatomic routing game, not all path-flow decompositions of an edge-flow at equilibrium are at equilibrium.*

Proof. Consider the graph on the right of Figure 2 with mean travel times for edges e_1, e_2, e_3, e_4 equal to $a, a+1, b, b-1$ for some $a, b > 0$, and standard deviations of travel times equal to $\sqrt{8}$, $\sqrt{3}$, 1, $\sqrt{8}$. The costs along the four possible paths are $Q_{13} = a+b+3$, $Q_{24} = a+b+3$, $Q_{14} = a+b+3$, and $Q_{23} = a+b+\sqrt{11}$, which are constants independent of the flow. The edge flow that sends $1/2$ unit of flow along each edge is an equilibrium, but only if decomposed properly. We know that an equilibrium can only use the minimum-cost paths π_{13} and π_{24}. (Although path π_{14} is also minimum-cost, using this path with the given edge flow would require sending flow on the higher-cost path π_{23} as well, to make the flow feasible, hence this cannot result in an equilibrium.) Viewed as a path-flow, the only decomposition that satisfies the equilibrium conditions is $f_{13} = f_{24} = 1/2$ and $f_{14} = f_{23} = 0$. Any other decomposition uses π_{23} and therefore does not provide an equilibrium. $\square$

A key insight from the above Lemma 6.1 is that *not all minimum-cost paths can be used in a decomposition*. Similarly, the social cost of different decompositions of a given edge-flow can vary.

Lemma 6.2. *In a nonatomic routing game, not all path-flow decompositions of an edge-flow minimize the social cost.*

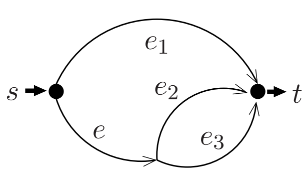

Player strategies	π_1	π_2	π_3
π_1	4.80, 4.80	4.80, 3.22	4.80, 3.11
π_2	3.22, 4.80	5.73, 5.73	4.73, 4.81
π_3	3.11, 4.80	4.81, 4.73	5.81, 5.81

FIGURE 3. Instance with no pure-strategy Nash equilibrium in an atomic unsplittable game with endogenous standard deviations. *Right:* Normal form game with two players.

Proof. Consider again the graph on the right of Figure 2 with mean travel times for edges e_1, e_2, e_3, e_4 equal to $2f_1, 6f_2, 4f_3, 10f_4$, and standard deviation of travel times equal to $0.6f_1, f_2, 2f_3, f_4$, respectively. Let the edge-flow be $0.3, 0.7, 0.4, 0.6$ respectively, and $\gamma = 1$. The costs along the four possible paths are $Q_{13} = 3.02$, $Q_{24} = 11.12$, $Q_{14} = 7.23$, and $Q_{23} = 6.86$, which are constants independent of the decomposition of the flow. Routing the unit demand at minimum cost, we get the flow $f_{13} = 0$, $f_{24} = 0.3$, $f_{14} = 0.3$, and $f_{23} = 0.4$, with a total cost of 8.24. The alternative of routing the flow as $f_{13} = 0.3$, $f_{24} = 0.6$, $f_{14} = 0$, and $f_{23} = 0.1$, provides a total cost of 8.27. □

The instance in the previous proof does not depend on the standard deviations being endogenous. An instance with exogenous standard deviations can be given where the structure does not change.

Remark 6.3. *A surprising fact raised by the example in Lemma 6.2 is that assigning flow to paths greedily does not provide a cost-minimizing flow. Indeed, the shortest path π_{13} carries zero flow.*

The lemmas above prompt the need of characterizing the structure of path-flow decompositions of equilibria and social optima. Does a succinct flow decomposition of an equilibrium or a social optimum always exist (namely one that assigns positive flows to only polynomially-many paths)? The following results answer this question in the positive. First, we prove that an edge-flow of a socially-optimal solution can be decomposed into a small number of paths.

Theorem 6.4. *For a social optimum $(x_e)_{e \in E}$ given as an edge-flow in the nonatomic case, there exists a* succinct *flow decomposition that uses at most $|E| + |K|$ paths.*

Proof. Because the edge-flow $\mathbf{x}$ is fixed, path costs are constant independent of the decomposition. Therefore, even though the cost functions $Q_\pi(\cdot)$ are nonlinear, the flow-decomposition problem can be written as the following linear program:

$$\begin{aligned} \min \sum_{\pi \in \mathcal{P}} & Q_\pi(\mathbf{x}) f_\pi \\ \text{s.t. } x_e = & \sum_{\pi \in \mathcal{P} \colon e \in \pi} f_\pi \quad \text{for } e \in E, \\ d_k = & \sum_{\pi \in \mathcal{P}_k} f_\pi \quad \text{for } k \in K, \\ f_\pi \geq & \, 0 \quad \text{for } \pi \in \mathcal{P}\,. \end{aligned} \tag{8}$$

The previous problem has $|E| + |K|$ equality constraints, from where the result follows because there is always an optimal solution to a linear program in which the number of non-zero variables is bounded by the number of equality constraints. □

Next, we prove a similar result for succinct equilibrium decompositions. The subtlety is that in the endogenous case, we do not even have an optimization formulation, let alone a linear programming formulation of the equilibrium as we do for the social optimum problem above. The insight is that for a *fixed* edge flow, the path costs are fixed. Therefore, equilibria of endogenous problems have corresponding equivalent equilibria of exogenous problems, in which the standard deviation values are *constant*.

Theorem 6.5. *For an equilibrium $(x_e)_{e \in E}$ given as an edge-flow of a nonatomic routing game, there exists a* succinct *flow decomposition that uses at most $|E| + |K|$ paths.*

Proof. For the exogenous setting, recall that the equilibrium is a solution to the convex program (3). For fixed edge flows, this is a linear program in the path-flow variables f_π with the same feasible set as the social optimum problem (8). Therefore, by the same argument as in Theorem 6.4, there exists a succinct flow decomposition that uses at most $|E| + |K|$ paths.

For the endogenous setting, we do not have an optimization formulation that can be turned into a linear program; however, since edge flows are fixed, so are the standard deviations corresponding to the edges. Therefore, an equilibrium in an endogenous setting has an equivalent equilibrium in the corresponding exogenous setting, in which the standard deviations are set to those in the endogenous formulation under the given edge flow. The edge flow must remain an equilibrium because costs along all paths do not change. Proceeding as in the previous paragraph, we obtain a succinct decomposition at equilibrium for the exogenous game, which will also be an equilibrium of the original formulation with endogenous standard deviations. □

It remains open whether finding an equilibrium path-flow decomposition from an equilibrium given as an edge flow can be done in polynomial time. This is related to the open question of whether the stochastic shortest path problem defined for finding a single minimum-cost path is in P [50].

7. Efficiency Analysis of Stochastic Wardrop Equilibria in Nonatomic Routing Games

In this section, we analyze the worst-case inefficiency of the mean-risk equilibria of the nonatomic routing game. To quantify this inefficiency, we make use of the concept of price of anarchy [54], first introduced by Koutsoupias and Papadimitriou [33] and used extensively in relation to transportation and telecommunications networks [61, 60, 15, 14, 55, 16].

The price of anarchy (POA) is defined as the supremum over all problem instances of the ratio of the equilibrium cost to the social optimum cost. A central planner would like to minimize $C(\mathbf{f})$ (see (2)) but typically a flow achieving that minimum is not possible when users make self-minded decisions and select paths that minimize their own costs, leading to an equilibrium outcome instead.

7.1. Exogenous Standard Deviations.

Prior research on deterministic routing games has shown that the price of anarchy is bounded by a relatively small constant. What is most surprising is that the inefficiency does not grow unbounded when networks become bigger and more complicated. Following Roughgarden [60], the price of anarchy is typically computed for a set of travel time functions given a-priori. We prove that in the case of stochastic travel times with exogenous standard deviations, the price of anarchy is the same as in the deterministic case. The bounds result from a modification of the bounding techniques of Correa *et al.* [15, 16].

For example, when the expected travel time functions are linear in the edge flow, the price of anarchy is $4/3$, meaning that, at equilibrium, the total cost experienced by users does not exceed 33.33% of that in an optimal solution. For nonlinear functions, such as those suggested by the Bureau of Public Roads [13], one needs to adjust the constant. When the expected travel times are degree-4 polynomials, as is typically used in practice, the price of anarchy evaluates to 2.151. We highlight that this is a worst-case bound, which by definition tends to be pessimistic. Although it is tight, because it was already tight for deterministic networks, the inefficiency of equilibria achieves it only in specially-constructed instances that are far from realistic. There is theoretical and computational research that tries to refine this to understand the worst-case inefficiency among 'realistic' instances [27, 32, 57, 16].

To prove the price-of-anarchy result in our stochastic setting, we use the same definition for the parameter β as Correa *et al.* [15]. Namely, we consider a family of expected travel time functions $\mathcal{L}$, and define for a travel time function $\ell \in \mathcal{L}$ and a number $v \geq 0$, $\beta(v, \ell) := \max_{x \geq 0}\{x(\ell(v) - \ell(x))\}/(v\ell(v))$, $\beta(\ell) := \sup_{v \geq 0} \beta(v, \ell)$, and finally $\beta(\mathcal{L}) := \sup_{\ell \in \mathcal{L}} \beta(\ell)$.

Theorem 7.1. *Consider a nonatomic routing game with continuous nondecreasing expected travel times belonging to a family $\mathcal{L}$ of travel time functions, and exogenous standard deviations. A stochastic Wardrop equilibrium $\mathbf{f}$ and a socially-optimal flow $\mathbf{f}^*$ minimizing the social cost (2) satisfy $C(\mathbf{f}) \leq (1-\beta(\mathcal{L}))^{-1}C(\mathbf{f}^*)$.*

Proof. We define the social cost of path-flow $\mathbf{x}$ under the prevailing path costs for the equilibrium $\mathbf{f}$ by $C^{\mathbf{f}}(\mathbf{x}) := Q(\mathbf{f}) \cdot \mathbf{x}$. Under this definition, $C(\mathbf{x}) = C^{\mathbf{x}}(\mathbf{x})$. Denote the (constant) path standard deviations by $\sigma_\pi := (\sum_{e \in \pi} \sigma_e^2)^{1/2}$. The variational inequality characterization of equilibria implies that $C(\mathbf{f}) \leq C^{\mathbf{f}}(\mathbf{x})$ for any feasible flow $\mathbf{x}$. Furthermore,

$$\begin{aligned} C^{\mathbf{f}}(\mathbf{x}) &= \sum_{e \in E} \ell_e(f_e)x_e + \sum_{\pi \in \mathcal{P}} \gamma\sigma_\pi x_\pi \leq \sum_{e \in E} \ell_e(x_e)x_e + \sum_{e \in E} \beta(\mathcal{L})\ell_e(f_e)f_e + \sum_{\pi \in \mathcal{P}} \gamma\sigma_\pi x_\pi \\ &= C(\mathbf{x}) + \beta(\mathcal{L}) \sum_{e \in E} \ell_e(f_e)f_e \leq C(\mathbf{x}) + \beta(\mathcal{L})C(\mathbf{f}). \end{aligned}$$

Here, the first inequality uses the definition of β and the second follows after completing the social cost function with the standard deviations. Therefore, $C(\mathbf{f}) \leq C(\mathbf{x})/(1-\beta(\mathcal{L}))$ for any feasible flow $\mathbf{x}$, implying that the price of anarchy is $(1 - \beta(\mathcal{L}))^{-1}$. $\square$

7.2. **Endogenous Standard Deviations.** In the case of endogenous standard deviations, an analysis of the price of anarchy is more elusive, not only for the complications of stochastic Wardrop equilibria but also because characterizing social optima in this case is difficult too. With the hope of simplifying the problem, we study the limiting case of extreme risk-aversion (the other 'extreme' case of risk-neutrality is already well-understood, as explained earlier). In that case, we consider that expected travel times are zero because users only care about standard deviations of travel time. Hence, 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 5.7 shows that there is a unique equilibrium that be computed efficiently with a convex program. We leave the case of intermediate values of risk aversion open.

For the extreme case of infinite risk aversion and polynomial standard deviation functions, we prove that the price of anarchy is one whenever the social cost function is convex. Unfortunately though, even in simple instances, the social cost is not convex because path-cost operators, although convex themselves, fail to be monotone as required. This, once more, happens because of the complicating square root. We finish by proving nevertheless that the price of anarchy is 1 but for instances with more restrictive assumptions.

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 functions are monomials of the same degree. Note that in the deterministic case, it is well known that the price of anarchy is exactly one precisely for monomials of the same degree [19].

Theorem 7.2. *Consider a nonatomic routing game with zero travel times and endogenous standard deviations of the form $\sigma_e(x_e) = a_e x_e^p$ for some fixed $p \geq 0$. A stochastic Wardrop equilibrium is a stationary point in the social-optimum (SO) problem that consists on minimizing $C(f)$ among feasible flows (see* (2)*).*

Proof. Consider the Lagrangian of the SO problem, $L(\mathbf{f}, \lambda) = \sum_{\pi\in\mathcal{P}} f_\pi Q_\pi(\mathbf{f}) + \lambda(1 - \sum_{\pi\in\mathcal{P}} f_\pi) - \sum_{\pi\in\mathcal{P}} \mu_\pi f_\pi$ (see, e.g., [10]). Its derivatives are

$$\frac{\partial L(\mathbf{f},\lambda)}{\partial f_\pi} = Q_\pi(\mathbf{f}) + \sum_{\pi'\in\mathcal{P}} f_{\pi'} \frac{\partial Q_{\pi'}(\mathbf{f})}{\partial f_\pi} - \lambda - \mu_\pi = Q_\pi(\mathbf{f}) + \sum_{\pi'\in\mathcal{P}} f_{\pi'} \frac{\sum_{e\in\pi\cap\pi'} \sigma_e(f_e)\sigma_e'(f_e)}{Q_{\pi'}(\mathbf{f})} - \lambda - \mu_\pi \,.$$

Let us evaluate the derivative above at the equilibrium flow $\mathbf{f}$ for a path π with flow $f_\pi > 0$. The multiplier μ_π can be discarded because the path carries positive flow and, hence, the corresponding constraint is not binding. Since the equilibrium conditions imply that path costs along flow-carrying paths are constant, we can replace the denominator $Q_{\pi'}(\mathbf{f})$ with $Q_\pi(\mathbf{f})$. Therefore, we have

$$\begin{aligned} Q_\pi(\mathbf{f}) + \frac{1}{Q_\pi(\mathbf{f})}\sum_{e\in\pi} \sigma_e(f_e)\sigma_e'(f_e)\Big(\sum_{\pi'\in\mathcal{P},\pi'\ni e} f_{\pi'}\Big) - \lambda &= Q_\pi(\mathbf{f}) + \frac{1}{Q_\pi(\mathbf{f})}\sum_{e\in\pi} \sigma_e(f_e)\sigma_e'(f_e) f_e - \lambda \\ = Q_\pi(\mathbf{f}) + \frac{p}{Q_\pi(\mathbf{f})}\sum_{e\in\pi} \sigma_e(f_e)^2 - \lambda &= Q_\pi(\mathbf{f}) + \frac{p}{Q_\pi(\mathbf{f})} Q_\pi(\mathbf{f})^2 - \lambda = (p+1)Q_\pi(\mathbf{f}) - \lambda. \end{aligned}$$

Here, we have used that for monomial standard deviation functions, $x_e\sigma_e'(x_e) = p a_e x_e^p = p\sigma_e(x_e)$. Setting $\lambda = (p+1)Q_\pi(\mathbf{f})$ results in $\partial L(\mathbf{f},\lambda)/\partial f_\pi = 0$ for paths π carrying positive flow at equilibrium. For a path π that does not carry flow, μ_π evaluates to a positive number because $Q_\pi(\mathbf{f}) \geq Q_{\pi'}(\mathbf{f})$ for any path π' that carries flow. Hence, the Kuhn-Tucker necessary conditions are satisfied at equilibrium. $\square$

As a corollary from the above theorem, whenever the SO problem has a unique stationary point (for example, if the social cost objective is strictly convex), 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 square root, the path costs are convex in the edge-flow variables when the standard deviations $\sigma_e(x_e)$ are convex functions.

Proposition 7.3. *The path costs $Q_\pi(\mathbf{x})$ are convex functions on $\mathbf{x}$ whenever the expected travel time and standard deviation functions are convex.*

Proof. Path costs $Q_\pi(\mathbf{x})$ can be split into two parts. The part corresponding to expected travel times on all the edges on π is convex because it is additive and the expected travel times on edges are convex. Hence, it suffices to show that the standard deviation component of the path cost is convex. Since the edge standard deviations $\sigma_e(x_e)$ are convex, we have that $\sigma_e(\beta x_e + (1-\beta)y_e) \leq \beta\sigma_e(x_e) + (1-\beta)\sigma_e(y_e)$ for all e, so it suffices to show that

$$\sqrt{\sum_{e\in\pi}[\beta\sigma_e(x_e) + (1-\beta)\sigma_e(y_e)]^2} \leq \beta\sqrt{\sum_{e\in\pi}\sigma_e(x_e)^2} + (1-\beta)\sqrt{\sum_{e\in\pi}\sigma_e(y_e)^2}\,,$$

for any two feasible flows $\mathbf{x}$ and $\mathbf{y}$, and $\beta \in [0,1]$. Squaring both sides and rearranging terms, the previous inequality is equivalent to

$$2\beta(1-\beta)\sum_{e\in\pi}\sigma_e(x_e)\sigma_e(y_e) \leq 2\beta(1-\beta)\sqrt{\Big[\sum_{e\in\pi}\sigma_e(x_e)^2\Big]\Big[\sum_{e\in\pi}\sigma_e(y_e)^2\Big]}\,.$$

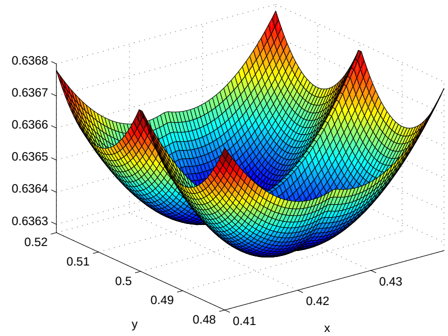

FIGURE 4. Non-convex slice of the social cost function corresponding to the instance used in Proposition A.1.

which holds if and only if $\left[\sum_{e\in\pi}\sigma_e(x_e)\sigma_e(y_e)\right]^2 \leq \left[\sum_{e\in\pi}\sigma_e(x_e)^2\right]\left[\sum_{e\in\pi}\sigma_e(y_e)^2\right]$. The last inequality is true because it is equivalent to $\sum_{e,i\in\pi}\left[\sigma_e(x_e)\sigma_i(y_i)-\sigma_i(x_i)\sigma_e(y_e)\right]^2 \geq 0$, proving the claim. □

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. The definition of a monotone operator is the same as a strictly-monotone one (6), except that the inequality is weak instead of strict.

Proposition 7.4. *The social cost $C(\mathbf{x})$ is convex whenever the path-cost operator Q is monotone and the path costs $Q_\pi(\mathbf{x})$ are convex.*

Proof. We need to prove that the social cost satisfies $C(\beta\mathbf{x}+(1-\beta)\mathbf{y}) \leq \beta C(\mathbf{x})+(1-\beta)C(\mathbf{y})$ for any two feasible flows $\mathbf{x}$ and $\mathbf{y}$, and $\beta\in[0,1]$. Using the convexity of path costs, it suffices to show that $[\beta\mathbf{x}+(1-\beta)\mathbf{y}]\,[\beta Q(\mathbf{x})+(1-\beta)Q(\mathbf{y})] \leq \beta\mathbf{x}Q(\mathbf{x})+(1-\beta)\mathbf{y}Q(\mathbf{y})$. This condition is equivalent to

$$\beta^2\mathbf{x}Q(\mathbf{x})+(1-\beta)^2\mathbf{y}Q(\mathbf{y})+\beta(1-\beta)\left[\mathbf{x}Q(\mathbf{y})+\mathbf{y}Q(\mathbf{x})\right] \leq \beta\mathbf{x}Q(\mathbf{x})+(1-\beta)\mathbf{y}Q(\mathbf{y})\,.$$

Regrouping the terms in one side and dividing over $\beta(1-\beta)$, we see that the last inequality holds because $(\mathbf{x}-\mathbf{y})\left[Q(\mathbf{x})-Q(\mathbf{y})\right] \geq 0$, by monotonicity of the path costs $Q(\cdot)$. □

Having convex path-cost functions may suggest that the social cost function is also convex. Unfortunately, the convexity of the latter fails to hold even in the basic case of linear standard deviation functions equal to $\sigma_e(x)=x$, as shown in Figure 4 and more formally proved in Proposition A.1 in the appendix. Nevertheless, we can still show that the POA is 1 in a network of n pairs of parallel edges connected in series (e.g., Figure 2(b) shows a network like this with 2 pairs of edges). Despite the limited class of topologies that are allowed, this example illustrates the difficulties one runs into when looking for bounds on the price of anarchy under the mean-risk objective. Although the equilibrium is characterized in the same way as when there is no uncertainty, analyzing optimal solutions is far from trivial, even with this specific structure. This should be contrasted to what was done for the case of exogenous standard deviations under general graphs and costs in Section 7.1. 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 travel time functions remains open.

Proposition 7.5. *Consider a nonatomic routing game on a network of n pairs of parallel edges connected in series. There is an end-to-end demand of 1 and the mean travel times are zero and the standard deviation functions are equal to $\sigma_e(x) = x$ in all edges e. For these instances, stochastic Wardrop equilibria and socially-optimal flows coincide.*

Proof. We take a flow that routes half a unit of demand in each edge, an decompose it arbitrarily in two paths. By symmetry, this must be an equilibrium. We want to prove that it also minimizes the social cost. To do that, for a given edge-flow, we characterize the decomposition that minimizes the social cost. Then, we will use the optimal decomposition to show that the edge-flow routes half a unit of demand in each edge, as the equilibrium.

It will be useful to refer to the set $\{i,\dots,n\}$ by $[i]$. For the instance we consider we can refer to a path by naming the top arcs that the path goes through. Indeed, we let π_X be the path that takes the top arcs in $X \subseteq [1]$ and the bottom arcs in $\bar{X}$ (the bar over X indicates the complement). We refer to the i-th pair of parallel edges as i for that on top and as $\bar{i}$ for that in the bottom.

Assume we are given an edge-flow $\mathbf{x}$ and we want to decompose it in a path-flow $(\mathbf{f}_\pi)_{\pi\in\mathcal{P}}$ to minimize its total cost $C(\mathbf{f}) = \sum_{\pi\in\mathcal{P}} \mathbf{f}_\pi Q_\pi(\mathbf{f})$, where $Q_\pi(\mathbf{f})$ can be assumed constant because it depends on the edge-flow $\mathbf{x}$ that is fixed. Without loss of generality, we can assume that $0 \le x_1 \le x_2 \le \dots \le x_n \le 1/2$ because we can reorder the the pairs and reverse the top and bottom edges.

Let us consider the following feasible decomposition of $\mathbf{x}$: $\mathbf{f}_{\pi_{[1]}} = x_1$, $\mathbf{f}_{\pi_{[i]}} = x_i - x_{i-1}$ for $i = 2,\dots,n$, $\mathbf{f}_{\pi_\emptyset} = 1 - x_n$, and zero for all other paths. We prove the optimality of the decomposition by induction. The basic case follows from the same argument as the inductive step. Assume that we are in step i of the induction. Then, paths $\pi_{[k]}$ for $k \in \{1,\dots,i-1\}$ are routed optimally and top arcs in $\{1,\dots,i-1\}$ are saturated. We will prove that path $\pi_{[i]}$ must be saturated, which happens when the flow on arc i reaches x_i, achieving the required decomposition.

If path $\pi_{[i]}$ does not carry flow because $x_{i-1} = x_i$, the path is already saturated, proving the inductive step. Otherwise, we have that $\pi_{[i]} > 0$. To prove the claim we will see that any feasible direction $(g_\pi)_{\pi\in\mathcal{P}}$ from $\mathbf{f}$ that removes flow from $\pi_{[i]}$ must induce a nonnegative gradient for the objective cost. The gradient is just $(Q_\pi(\mathbf{f}))_{\pi\in\mathcal{P}}$ because the path costs do not depend on the decomposition. A feasible direction in this problem is a path-flow in $\mathbb{R}^{\mathcal{P}}$ whose corresponding edge-flow is zero for all arcs. The circulation may have a negative flow along a path π, but only if the $\mathbf{f}_\pi > 0$ because one cannot remove flow from a path that does not carry it.

We see this first for simple directions and then note that an arbitrary feasible direction can be decomposed in a conic combination of these simple directions. For $W \subseteq [i]$, the simple direction associated with W consists on removing flow from the path that takes all top arcs and all bottom arcs from stage i onwards and adding that flow in two complementary paths that take the top arcs in W and the bottom arcs in W, and vice-versa. Concretely, $g^W_{\pi_{[i]}} = g^W_{\pi_\emptyset} = -1$ and $g^W_{\pi_W} = g^W_{\pi_{[i]\setminus W}} = 1$. It is straightforward to see that g^W is a feasible direction because the resulting edge-flow is zero and $\mathbf{f}^W_{\pi_{[i]}} > 0$ by assumption and $\mathbf{f}^W_{\pi_\emptyset} > 0$ because $x_n < 1/2$. We must prove that the gradient along g^W is not an improving direction, $Q_{\pi_W} + Q_{\pi_{[i]\setminus W}} - Q_{\pi_{[i]}} - Q_{\pi_\emptyset} \ge 0$. By grouping the terms appropriately and letting $B_0 = \sum_{k\in\{1\dots,i-1\}} x^2_{\bar{k}}$, $T_W = \sum_{k\in W} x^2_k$, $T_{\bar{W}} = \sum_{k\in[i]\setminus W} x^2_k$, $B_W = \sum_{k\in W} x^2_{\bar{k}}$, and $B_{\bar{W}} = \sum_{k\in[i]\setminus W} x^2_{\bar{k}}$, the previous inequality can be written as

$$\sqrt{B_0 + T_W + B_{\bar{W}}} + \sqrt{B_0 + T_{\bar{W}} + B_W} \geq \sqrt{B_0 + T_W + T_{\bar{W}}} + \sqrt{B_0 + B_W + B_{\bar{W}}}\,.$$

Squaring both sides, canceling equal terms and squaring again, we get that the previous is equivalent to

$$(B_0 + T_W + B_{\bar{W}})(B_0 + T_{\bar{W}} + B_W) \geq (B_0 + T_W + T_{\bar{W}})(B_0 + B_W + B_{\bar{W}})\,.$$

Distributing and simplifying, we get $(T_W - B_W)(T_{\bar{W}} - B_{\bar{W}}) \geq 0$, which must hold because both terms are negative since $x_k \leq 1/2 \leq x_{\bar{k}}$ for all k.

By linear algebra arguments, we can see that an arbitrary feasible direction g can be expressed as $\sum_{W \subseteq [i]} \lambda_W g^W$ for $\lambda_W \geq 0$, therefore proving that the gradient along g is also nonnegative. Because decreasing the flow on path $\pi_{[i]}$ cannot improve the objective, $\pi_{[i]}$ must be saturated in an optimal decomposition. The rest of the flow can only take the bottom arc $\bar{i}$, completing the inductive step. Therefore, the minimal social cost as a function of the edge flow $\mathbf{x}$ satisfies:

$$\begin{aligned}
C(x_1, \ldots, x_n) \quad &= \quad x_1\sqrt{x_1^2 + x_2^2 + \ldots + x_n^2} + (x_2 - x_1)\sqrt{(1-x_1)^2 + x_2^2 + \ldots + x_n^2} \\
&\quad\quad + \ldots + (1 - x_n)\sqrt{(1-x_1)^2 + \ldots + (1-x_n)^2} \\
&\geq \quad \frac{1}{\sqrt{n}}[x_1\,(x_1 + x_2 + \ldots + x_n) + (x_2 - x_1)\,((1-x_1) + x_2 + \ldots + x_n) \\
&\quad\quad + \ldots + (1-x_n)\,((1-x_1) + \ldots + (1-x_n))] \\
&= \quad \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left[x_i^2 + (1-x_i)^2\right] \geq \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\frac{1}{2} = \frac{\sqrt{n}}{2},
\end{aligned}$$

where the first inequality follows from applying the root-mean square inequality $a_1^2 + \ldots + a_n^2 \geq (a_1 + \ldots + a_n)^2/n$ to every square root term. Therefore, $C(x_1, \ldots, x_n) \geq C(1/2, \ldots, 1/2)$. Since the cost of the equilibrium is a lower bound for the cost of an arbitrary flow, the claim holds. □

8. Conclusions and Open Problems

We have set out to extend the classical theory of Wardrop equilibria and network assignment problems to the more realistic setting of uncertain travel times. In this work, we have focused exclusively on theoretical questions about the nature of the competition. The uncertainty of travel times calls for models that incorporate users' attitudes towards risk, which we have captured through a linear combination of the expectation and the standard deviation of travel times along the chosen route.

We have considered nonatomic and atomic routing games, each with exogenous or endogenous standard deviation functions and provided results on (1) the existence and characterization of equilibria; (2) succinct path decompositions of equilibrium and socially optimal flows; (3) the inefficiency of equilibria. The directions pursued in this work have opened many other questions that 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?

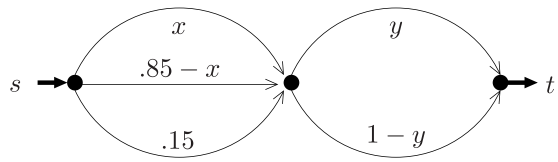

FIGURE 5. Instance that illustrates a non-convex social cost function.

- 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 users with heterogenous attitudes toward risk [53]. Can some of the results in this paper be extended to that setting?

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

Acknowledgements. Most of the work was done while the first author was at MIT CSAIL. We thank Hari Balakrishnan, José Correa, Costis Daskalakis and Asu Ozdaglar for helpful discussions. We also thank Darrell Hoy for carefully proofreading this manuscript and for suggestions on improving the presentation. The work of the first author was supported by the National Science Foundation grant number 0931550. The conference version of this paper appeared in SAGT'11 [51]

APPENDIX A. ADDITIONAL PROOFS

The following theorem shows that the convexity of the social cost fails to hold even in the basic case of linear standard deviation functions equal to $\sigma_e(x) = x$.

Proposition A.1. *The social cost* $\sum_{\pi \in \mathcal{P}} x_\pi (\sum_{e \in \pi} x_e^2)^{1/2}$ *is not convex.*

Proof. Consider the network of 3 parallel edges, followed by 2 parallel edges shown in Figure 5. Taking the same flows as in the pseudo-monotonicity counterexample given in Proposition 5.6 violates the convexity condition

$$Q(\beta \mathbf{x} + (1-\beta)\mathbf{y}) \leq \beta Q(\mathbf{x}) + (1-\beta) Q(\mathbf{y})$$

for $\beta = 1/2$. In particular, we denote a path flow vector by $\mathbf{x} = [x_{14}, x_{24}, x_{34}, x_{15}, x_{25}, x_{35}]$ where the subscript ij denotes the path using edges e_i, e_j with the notation from the figure. Taking flows $\mathbf{x} = [0, 0, 0.1, 0.2, 0.7, 0]$ and $\mathbf{y} = [0.1, 0, 0, 0, 0.7, 0.2]$ yields corresponding path costs

$$\begin{aligned}
Q(\mathbf{x}) &= [\sqrt{0.05}, \sqrt{0.5}, \sqrt{0.02}, \sqrt{0.85}, \sqrt{1.3}, \sqrt{0.82}] \\
Q(\mathbf{y}) &= [\sqrt{0.02}, \sqrt{0.5}, \sqrt{0.05}, \sqrt{0.82}, \sqrt{1.3}, \sqrt{0.85}] \\
Q\Big(\frac{\mathbf{x}+\mathbf{y}}{2}\Big) &= [\sqrt{0.0325}, \sqrt{0.5}, \sqrt{0.0325}, \sqrt{0.8325}, \sqrt{1.3}, \sqrt{0.8325}].
\end{aligned}$$

After some algebra, we have a counterexample to the convexity of the social cost function:

$$\frac{\mathbf{x}+\mathbf{y}}{2} \cdot Q\Big(\frac{\mathbf{x}+\mathbf{y}}{2}\Big) = 0.99863 > 0.99666 = \frac{1}{2}\mathbf{x} \cdot Q(\mathbf{x}) + \frac{1}{2}\mathbf{y} \cdot Q(\mathbf{y}).$$

□

The question of the uniqueness of equilibrium in the nonatomic setting with endogenous standard deviations invites the seemingly simpler question of whether, for fixed path cost values for all paths, there exists a unique edge flow vector that attains them. We have not seen this question addressed even for the deterministic setting. Below, we show that it is true for the latter for linear travel times but leave open whether this statement is true in the stochastic setting.

Lemma A.2. *Consider a deterministic nonatomic routing game with linear travel time functions. Suppose we are given a (valid) vector of path cost values $\{c_\pi\}_{\pi\in\mathcal{P}}$, for all possible paths. Then, there is a unique edge-flow vector that attains those costs.*

Proof. Denote by Δ the feasible edge flow polytope in $\mathbb{R}^{|E|}$. We define the following auxiliary optimization problem:

$$\min_{\mathbf{f}\in\Delta} H(\mathbf{f}) := \sum_{\pi\in\mathcal{P}} \left(\sum_{e\in\pi} \ell_e(f_e) - c_\pi \right)^2 . \tag{9}$$

Clearly, the minimum of $H(\mathbf{f})$ is zero, which is attained when the path costs are given by the flow that generated the cost values c_π. To show that (9) has a unique solution, it suffices to show that the objective function is strictly convex, which can be done by proving that the gradient of the objective is strictly monotone. Hence, we need to prove that

$$(\mathbf{f}-\mathbf{f}')\cdot\big(\nabla H(\mathbf{f}) - \nabla H(\mathbf{f}')\big) > 0, \qquad \forall \mathbf{f},\mathbf{f}' \text{ with } \mathbf{f}\neq\mathbf{f}'. \tag{10}$$

Denote the travel time functions by $\ell_i(f_i) = a_i f_i + b_i$, for all edges $i\in E$. Then, $\partial\ell_i(f_i)/\partial f_j = a_i\mathcal{I}_{i=j}$, where $\mathcal{I}$ is the indicator function and

$$\frac{\partial H(\mathbf{f})}{\partial f_i} = \sum_{\pi\in\mathcal{P}} 2\Big(\sum_{j\in\pi}\ell_j(f_j) - c_\pi\Big)\frac{\partial \ell_j(f_j)}{\partial f_i} .$$

The LHS of (10) becomes

$$\begin{aligned}
\sum_{i\in E}(f_i - f_i')\left(\frac{\partial H(\mathbf{f})}{\partial f_i} - \frac{\partial H(\mathbf{f}')}{\partial f_i}\right) &= \sum_{i\in E}(f_i-f_i')\Big[\sum_{\pi\ni i}2a_i(\sum_{j\in\pi}\ell_j(f_j)-c_\pi) - \sum_{\pi\ni i}2a_i(\sum_{j\in\pi}\ell_j(f_j')-c_\pi)\Big] \\
&= \sum_{i\in E}(f_i-f_i')\left[\sum_{\pi\ni i}2a_i\sum_{j\in\pi}(\ell_j(f_j)-\ell_j(f_j'))\right] \\
&= \sum_{i\in E}(f_i-f_i')\left[\sum_{\pi\ni i}2a_i\sum_{j\in\pi}a_j(f_j-f_j')\right] \\
&= \sum_{i\in E}\sum_{\pi\ni i}2a_i(f_i-f_i')\left[\sum_{j\in\pi}a_j(f_j-f_j')\right] \\
&= \sum_{\pi\in\mathcal{P}}\sum_{i\in\pi}2a_i(f_i-f_i')\left[\sum_{j\in\pi}a_j(f_j-f_j')\right]
\end{aligned}$$

$$= 2 \sum_{\pi \in \mathcal{P}} \left[\sum_{i \in \pi} a_i (f_i - f_i') \right]^2 \geq 0.$$

In particular, we see that the above expression is strictly positive for all pairs of flow vectors that are not equal, $\mathbf{f} \neq \mathbf{f}'$. $\square$

The proof exploits the fact that the travel time functions are linear in at least two places: first, where the travel time derivatives simplify to constants and second, where the difference of travel times for two different flow values $\ell_e(f_e) - \ell_e(f_e')$ can be represented as a multiple of the difference between these flow values $(f_e - f_e')$. Consequently, the cross products of flow value differences of pairs of edges, whose sign is not necessarily positive, can be conveniently subsumed in the square of the sum of differences of flow values over edges so that the overall expression becomes positive. Thus, it is not immediately clear if such a proof is possible for nonlinear travel time functions. In fact, it is not even clear if uniqueness holds for general travel time functions.

References

[1] H. Z. Aashtiani and T. L. Magnanti. Equilibria on a congested transportation network. *SIAM Journal on Algebraic and Discrete Methods*, 2(3):213–226, 1981.

[2] T. Akamatsu. Cyclic flows, Markov processes and stochastic traffic assignment. *Transportation Research*, 30B(5):369–386, 1996.

[3] 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.

[4] G. Andreatta and L. Romeo. Stochastic shortest paths with recourse. *Networks*, 18:193–204, 1988.

[5] 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.

[6] J.-B. Baillon and R. Cominetti. Markovian traffic equilibrium. *Mathematical Programming*, 111B:33–56, 2008.

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

[8] M. Bell. Alternatives to Dial's logit assignment algorithm. *Transportation Research*, 29B:287–295, 1995.

[9] 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.

[10] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, MA, 1999.

[11] D. P. Bertsekas and J. N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.

[12] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98(1-3):49–71, 2003.

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

[14] C. K. Chau and K. M. Sim. The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. *Operations Research Letters*, 31(5):327–334, 2003.

[15] 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.

[16] 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.

[17] J. R. Correa and N. E. Stier-Moses. Wardrop equilibria. In J. J. Cochran, editor, *Encyclopedia of Operations Research and Management Science*. Wiley, 2011.

[18] S. C. Dafermos. Traffic equilibrium and variational inequalities. *Transportation Science*, 14(1):42–54, 1980.

[19] 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.

[20] C. Daganzo and Y. Sheffi. On stochastic models of traffic assignment. *Transportation Science*, 11(3):253–274, 1977.

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

[22] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:S. 269271, 1959.

[23] Y. Y. Fan, R. E. Kalaba, and J. E. Moore. Arriving on time. *Journal of Optimization Theory and Applications*, 127(3):497–513, 2005.

[24] 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.

[25] C. Fisk. Some developments in equilibrium traffic assignment. *Transportation Research*, 14B(3):243–255, 1980.

[26] H. Föllmer and A. Schied. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, 2004.

[27] E. J. Friedman. Genericity and congestion control in selfish routing. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, volume 5, pages 4667–4672, Atlantis, Paradise Island, Bahamas, 2004.

[28] S. Gabriel and D. Bernstein. The traffic equilibrium problem with nonadditive path costs. *Transportation Science*, 31:337–348, 1997.

[29] P. T. Harker. Multiple equilibrium behaviors of networks. *Transportation Science*, 22(1):39–46, 1988.

[30] P. T. Harker and J.-S. Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Math. Program.*, 48:161–220, September 1990.

[31] P. Hartman and G. Stampacchia. On some nonlinear elliptic differential functional equations. *Acta Math.*, 115:153–188, 1966.

[32] O. Jahn, R. H. Möhring, A. S. Schulz, and N. E. Stier-Moses. System-optimal routing of traffic flows with user constraints in networks with congestion. *Operations Research*, 53(4):600–616, 2005.

[33] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. Conference version in *16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, Trier, Germany, 1999.

[34] P. Krokhmal, M. Zabarankin, and S. Uryasev. Modeling and optimization of risk. *Surveys in Operations Research and Management Science*, 16(2):49–66, 2011.

[35] G. Langer. Traffic in the united states. *ABC News*, February 13 2005.

[36] T. Larsson, Z. Liu, and M. Patriksson. A dual scheme for traffic assignment problems. *Optimization*, 42:323–358, 1997.

[37] 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.

[38] H. K. Lo and Y.-K. Tung. Network with degradable links: capacity analysis and design. *Transportation Research*, 37B(4):345–363, 2003.

[39] M. Maher. Algorithms for logit-based stochastic user equilibrium assignment. *Transportation Research*, 32B(8):539–549, 1998.

[40] H. M. Markowitz. *Mean-Variance Analysis in Portfolio Choice and Capital Markets*. Basil Blackwell, Cambridge, MA, 1987.

[41] P. B. Mirchandani and H. Soroush. Generalized traffic equilibrium with probabilistic travel times and perceptions. *Transportation Science*, 21:133–152, 1987.

[42] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[43] J. F. Nash. Noncooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

[44] Y. Nie. Multi-class percentile user equilibrium with flow-dependent stochasticity. *Transportation Research Part B*, 45:1641–1659, 2011.

[45] Y. Nie and X. Wu. Reliable a priori shortest path problem with limited spatial and temporal dependencies. In *Proceedings of the 18th International Symposium on Transportation and Traffic Theory*, pages 169–196, 2009.

[46] Y. Nie and X. Wu. Shortest path problem considering on-time arrival probability. *Transportation Research Part B*, 43:597–613, 2009.

[47] E. Nikolova. *Strategic algorithms*. PhD thesis, Massachusetts Institute of Technology. Dept. of Electrical Engineering and Computer Science, 2009.

[48] 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.

[49] 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.

[50] 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.

[51] E. Nikolova and N. E. Stier-Moses. Stochastic selfish routing. In *SAGT*, 2011.

[52] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK, 2007.

[53] F. Ordóñez and N. E. Stier-Moses. Wardrop equilibria with risk-averse users. *Transportation Science*, 44(1):63–86, 2010.

[54] C. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, STOC '01, pages 749–753, New York, NY, USA, 2001. ACM.

[55] G. Perakis. The "price of anarchy" under nonlinear and asymmetric costs. *Mathematics of Operations Research*, 32(3):614–628, 2007.

[56] A. C. Pigou. *The Economics of Welfare*. Macmillan, London, 1920.

[57] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker. On selfish routing in Internet-like environments. *Transactions on Networking*, 14(4):725–738, 2006.

[58] T. Rockafellar. Coherent approaches to risk in optimization under uncertainty. *Tutorials in operations research*, pages 38–61, 2007.

[59] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[60] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[61] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[62] W. H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97(1):81 – 108, 2001.

[63] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2011.

[64] D. Schrank, T. Lomax, and S. Turner. *Annual Urban Mobility Report*. Texas Transportation Institute, 2010. Available at `http://mobility.tamu.edu/ums`.

[65] Y. Sheffi. *Urban Transportation Networks*. Prentice-Hall, Englewood, NJ, 1985.

[66] Y. Sheffi and W. Powell. An algorithm for the traffic assignment problem with random link costs. *Networks*, 12:191–207, 1982.

[67] M. J. Smith. The existence, uniqueness and stability of traffic equilibria. *Transportation Research*, 13B(4):295–304, 1979.

[68] M. Trahan. Probabilistic assignment: An algorithm. *Transportation Science*, 8(4):311–320, 1974.

[69] T. Uchida and Y. Iida. Risk assignment: a new traffic assignment model considering risk of travel time variation. In C. F. Daganzo, editor, *Proceedings of the 12th International Symposium on Transportation and Traffic Theory*, pages 89–105. Elsevier, Amsterdam, 1993.

[70] S. V. Ukkusuri and S. T. Waller. Approximate analytical expressions for transportation network performance under demand uncertainty. *Transportation Letters: The International Journal of Transportation Research*, 2(2):111–123, 2010.

[71] 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.

[72] X. Wu and Y. Nie. Modeling heterogeneous risk-taking behavior in route choice: A stochastic dominance approach. *Transportation Research Part A*, 45:896–915, 2011.