

Robust Adaptive Routing Under Uncertainty

Arthur Flajolet

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, flajolet@mit.edu
Ecole Polytechnique, Route de Saclay, 91120 Palaiseau, France

Sébastien Blandin

IBM Research Collaboratory, 9 Changi Business Park Central 1, Singapore 486048, Singapore, sblandin@sg.ibm.com

Patrick Jaillet

Department of Electrical Engineering and Computer Science, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, jaillet@mit.edu

We consider the problem of finding an optimal history-dependent routing strategy on a directed graph weighted by stochastic arc costs when the objective is to minimize the risk of spending more than a prescribed budget. To help mitigate the impact of the lack of information on the arc cost probability distributions, we introduce a robust counterpart where the distributions are only known through confidence intervals on some statistics such as the mean, the mean absolute deviation, and any quantile. Leveraging recent results in distributionally robust optimization, we develop a general-purpose algorithm to compute an approximate optimal strategy. To illustrate the benefits of the robust approach, we run numerical experiments with field data from the Singapore road network.

Key words: stochastic shortest path; Markov decision process; robust optimization

1. Introduction

1.1. Motivation

Stochastic Shortest Path (SSP) problems have emerged as natural extensions to the classical shortest path problem when arc costs are uncertain and modeled as outcomes of random variables. In particular, we consider in this paper the class of adaptive SSPs, which can be formulated as *Markov Decision Processes* (MDPs), where we optimize over all history-dependent strategies. As standard with MDPs, optimal policies are characterized by dynamic programming equations involving expected values (e.g. Bertsekas and Tsitsiklis (1991)). Yet, computing the expected value of a function of a random variable generally requires a full description of its probability distribution, and this can be hard to obtain accurately due to errors and sparsity of measurements. In practice, only finite samples are available and an optimal strategy based on approximated arc cost probability distributions may be suboptimal with respect to the real arc cost probability distributions.

One of the most common applications of SSPs deals with the problem of routing vehicles in transportation networks. Providing driving itineraries is a challenging task as suppliers have to cope simultaneously with limited knowledge about random fluctuations in traffic congestion (e.g. caused

by traffic incidents, variability of travel demand) and users' desire to arrive on time. These considerations have led to the definition of the Stochastic On-Time Arrival (SOTA) problem, an adaptive SSP problem with the objective of maximizing the probability of on-time arrival, and formulated using dynamic programming in Nie and Fan (2006). The algorithm proposed in Samaranayake et al. (2012b) to solve this problem assumes the knowledge of the complete arc travel-time distributions. Yet, in practice, such distributions tend to be estimated from samples which are sparse and error-prone.

In recent years, *Distributionally Robust Optimization* (DRO) has emerged as a new framework for decision-making under uncertainty when the underlying distributions are only known through some statistics or from collections of samples. DRO was put forth in an effort to capture both risk (uncertainty on the outcomes) and ambiguity (uncertainty on the probabilities of the outcomes) when optimizing over a set of alternatives, thus lying at the crossroad between stochastic and robust optimization. The computational complexity of this approach can vary greatly, depending on the nature of the ambiguity sets and on the structure of the optimization problem, see Wiesemann et al. (2014) and Delage and Ye (2010) for convex problems, and Calafiore and Ghaoui (2006) for chance-constraint problems. Even in the absence of decision variables, the theory proves useful in order to derive either numerical or closed form bounds on expected values using tools drawn from linear programming, e.g. tailored dual simplex algorithm in Prékopa (1990), and from semidefinite programming as in Bertsimas and Popescu (2005) and Vandenberghe et al. (2007).

In the case of limited knowledge of the arc cost probability distributions, we propose to bring DRO to bear on adaptive SSP problems to help mitigate the impact of the lack of information and introduce Distributionally Robust Adaptive Stochastic Shortest Path problems. Our work fits into the literature on distributionally robust MDPs where the transition probabilities are only known to lie in prescribed ambiguity sets (e.g. Nilim and Ghaoui (2005), Iyengar (2005), Xu and Mannor (2010), and Wiesemann et al. (2013)). While some of the methods developed in the aforementioned literature can be shown to carry over, adaptive SSPs exhibit a particular structure that allows for a large variety of ambiguity sets and enables the development of faster solution procedures. Specifically, optimal strategies for finite-horizon distributionally robust MDPs are characterized by a Bellman recursion on the worst-case expected reward-to-go. While standard approaches focus on computing this last quantity for each state independently from one another, closely related problems (e.g. estimating an expected value $\mathbb{E}[f(t - X)]$ where the random variable X is fixed but t varies depending on the state) carry across states for adaptive SSPs, and, as a result, making the most of previous computations becomes crucial to achieve computational tractability.

1.2. Related Work

Extending the shortest path problem by assigning random, as opposed to deterministic, costs to arcs requires some additional modeling assumptions. Over the years, many formulations have been proposed which differ along three main features:

- The specific objective function to optimize: in the presence of uncertainty, the most natural approach is to minimize the total expected costs, see Bertsekas and Tsitsiklis (1991), and Miller-Hooks and Mahmassani (2000) for time-dependent random costs. However, this approach is oblivious to risk. In an attempt to take that factor into account, Loui (1983) proposed earlier to rely on utility functions of moments (e.g. mean costs and variances) involving an inherent trade-off, and considered multi-objective criteria. However, Bellman’s principle of optimality no longer holds for arcs weighted by multidimensional costs, giving rise to computational hardness. A different approach consists of introducing a budget, set by the user, corresponding to the maximum total cost he is willing to pay to reach his terminal node. Such approaches have been considered along several different directions. Research efforts have considered either minimizing the probability of budget overrun (see Frank (1969), Nikolova et al. (2006b), and also Xu and Mannor (2011) for probabilistic goal MDPs), minimizing more general functions of the budget overrun as in Nikolova et al. (2006a), minimizing refined satisficing measures in order to guarantee good performances with respect to several other objectives as in Jaillet et al. (2015), and constraining the probability of over-spending while optimizing the expected costs as in Xu et al. (2012).

- The admissible set of strategies over which we are free to optimize: incorporating uncertainty may cause history-dependent strategies to significantly outperform a priori paths depending on the chosen performance index. This is the case for the SOTA problem where two types of formulations have been considered: (i) an a-priori formulation which consists in finding a path before taking any actions, see Nikolova et al. (2006b) and Nie and Wu (2009); and (ii) an adaptive formulation which allows to update the path to go based on the remaining budget, see Nie and Fan (2006), Samaranayake et al. (2012b), and Parmentier and Meunier (2014).

- The knowledge on the random arc costs taken as an input: it can range from the full knowledge of the probability distributions to having access to only a few samples drawn from them. In practical settings, the problem of estimating accurately some statistics (e.g. mean cost and variance) seems more reasonable than retrieving the full probability distribution. For instance, Jaillet et al. (2015) consider lower-order statistics (minimum, average and maximum costs) and make use of closed form bounds derived in the DRO theory. These considerations were extensively investigated in the context of distributionally robust MDPs. From a theoretical standpoint, Wiesemann et al. (2013) show that a property coined as rectangularity has to be satisfied by the ambiguity sets for computational

tractability, while Iyengar (2005) characterizes the optimal policies with a dynamic programming equation for general rectangular ambiguity sets. The ambiguity sets are parametric in Wiesemann et al. (2013), where the parameter lies in the intersection of finitely many ellipsoids, are based on likelihood measures in Nilim and Ghaoui (2005), and are defined by linear inequalities in White III and Eldeib (1994).

We give an overview of prior formulations in Table 1.

Table 1 Literature review.

Author(s)	Objective function	Strategy	Uncertainty description	Approach
Loui (1983)	utility function	a priori	moments	dominated paths
Nikolova et al. (2006b)	probability of budget overrun	a priori	normal distributions	convex optimization
Nie and Fan (2006) Samaranayake et al. (2012b)	probability of budget overrun	adaptive	distributions	dynamic programming
Nilim and Ghaoui (2005)	expected cost	adaptive	maximum-likelihood ambiguity sets	dynamic programming
Jaillet et al. (2015) Adulyasak and Jaillet (2014)	requirements violation	a priori	distributions or moments	iterative procedure
Gabrel et al. (2013)	worst-case cost	a priori	intervals or discrete scenarios	integer programming
Parmentier and Meunier (2014)	monotone risk measure	a priori	distributions	labeling algorithm
Our work	risk function of the budget overrun	adaptive	distributions or confidence intervals on statistics	dynamic programming

1.3. Contributions

The main contributions of this paper can be summarized as follows:

1. We extend the class of adaptive SSP problems, first introduced in Fan et al. (2005) when the objective is to minimize the probability of budget overrun, to general risk functions of the budget overrun. We characterize optimal strategies and identify conditions on the risk function under which infinite cycling is provably suboptimal. For any risk function satisfying these conditions, we provide an efficient solution procedure to compute an ϵ -approximate optimal strategy for any $\epsilon > 0$.
2. We introduce the distributionally robust version of this general problem, under rectangular ambiguity sets. We characterize optimal robust strategies and extend the conditions ruling out infinite cycling. For any risk function satisfying these conditions, we provide efficient solution

procedures to compute an ϵ -approximate optimal strategy when the arc cost distributions are only known through confidence intervals on piecewise affine statistics (e.g. the mean, the mean absolute deviation, any quantile...) for any $\epsilon > 0$.

Special cases where the objective is to minimize the probability of budget overrun and the arc costs are independent and take on values that are multiple of a unit cost can serve as a basis for comparison with prior work on distributionally robust MDPs. For this subclass of problems, our formulation can be interpreted as a distributionally robust MDP with finite horizon N , finitely many states n (resp. actions m), and a rectangular ambiguity set. Using the solution methodology developed in this paper, we can compute an ϵ -optimal strategy with complexity $O(m \cdot n \cdot \log(\frac{N}{\epsilon}) \cdot \log(n))$.

The remainder of the paper is organized as follows. In Section 2, we introduce the adaptive SSP problem and its distributionally robust counterpart. Section 3 (resp. Section 4) is devoted to the theoretical and computational analysis of the nominal (resp. robust) problem. In Section 5, we consider a vehicle routing application and present results of numerical experiments run with field data from the Singapore road network. In Section 6, we relax some of the assumptions made in Section 2 and extend the results presented in Sections 3 and 4.

Notations For a function $g(\cdot)$ and a random variable X distributed according to p , we denote by $\mathbb{E}_{X \sim p}[g(X)]$ the expected value of $g(X)$. For a set $S \subset \mathbb{R}^n$, \bar{S} is the closure of S in the standard topology of \mathbb{R}^n , $\text{conv}(S)$ denotes the convex hull generated by S and $|S|$ denotes the cardinality of S . For a set $S \subset \mathbb{R}^2$, \hat{S} denotes the upper convex hull of S , i.e. $\hat{S} = \{(x, y) \in \mathbb{R}^2 : \exists (a, b) \in \text{conv}(S) \text{ such that } x = a \text{ and } y \geq b\}$.

2. Problem Formulation

In this section, we formulate the adaptive SSP problem for any risk function of the budget overrun. Then, we introduce the distributionally robust approach based on ambiguity sets to tackle the situation of limited knowledge of the arc cost probability distributions.

2.1. Nominal problem

Let $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ be a finite directed graph where each arc $(i, j) \in \mathcal{A}$ is assigned a collection of non-negative random costs $(c_{ij}^\tau)_{\tau \geq 0}$. We consider a user traveling through \mathcal{G} leaving from s and wishing to reach d within a total prescribed budget T . Having already spent a total cost τ and being at node i , choosing to cross arc (i, j) would incur an additional cost c_{ij}^τ , whose value becomes known after the arc is crossed. In vehicle routing applications, c_{ij}^τ typically models the travel time along arc (i, j) at time τ and T is the deadline imposed at the destination. The objective is to find a strategy to reach d maximizing a risk function of the budget overrun, denoted by $f(\cdot)$. Mathematically, this corresponds to solving:

$$\sup_{\pi \in \Pi} \mathbb{E}[f(T - X_\pi)], \quad (1)$$

where Π is the set of all history-dependent randomized strategies, i.e. mappings from the past realizations of the costs and the previously visited nodes to probability distributions over the set of neighboring nodes, and X_π is the random cost associated with strategy π when leaving from node s with budget T . We denote by \mathcal{H} the set of all possible histories of the previously experienced costs and previously visited nodes. Examples of natural risk functions include $f(t) = t \cdot 1_{t \leq 0}$, $f(t) = 1_{t \geq 0}$, and $f(t) = -|t|$ which translate into, respectively, minimizing the expected budget overrun, maximizing the probability of completion within budget, and penalizing the expected deviation from the target budget. When $f(t) = 1_{t \geq 0}$, we recover the adaptive SOTA problem introduced in Fan et al. (2005). We will restrict our attention to risk functions satisfying natural properties meant to prevent infinite cycling in Theorem 1 of Section 3.1, e.g. maximizing the expected budget overrun is not allowed. Without any additional assumption on the random costs, (1) is computationally intractable and characterizing an optimal solution is theoretically hard. To simplify the problem, a common approach in the literature is to assume independence of the arc costs, see for example Fan et al. (2005) and Jaillet et al. (2015).

ASSUMPTION 1. $(c_{ij}^\tau)_{(i,j) \in \mathcal{A}, \tau \geq 0}$ are independent random variables.

In practice, the costs of neighboring arcs can be highly correlated for some applications and Assumption 1 may then appear unreasonable. It turns out that most of the results derived in this paper can be extended to the case where the dependence can be modeled by *Markov* chains of finite order, i.e., where the cost of an arc depends on the past $m \in \mathbb{N}$ experienced costs. This is of course at the price of more technicalities and an increased complexity both in terms of modeling and of computational requirements. To simplify the presentation, Assumption 1 is used throughout most of the paper and the extension to *Markov* chains is discussed in Section 6.1. For the same reason, we further assume that the random costs are identically distributed across τ .

ASSUMPTION 2. For each arc $(i, j) \in \mathcal{A}$, the distribution of c_{ij}^τ does not depend on τ .

The extension to τ -dependent arc cost distributions is detailed in Section 6.2. For clarity of the exposition, we omit the superscript τ in the notations when it is unnecessary and simply denote the costs by $(c_{ij})_{(i,j) \in \mathcal{A}}$, even though the cost of an arc corresponds to an independent realization of its corresponding random variable each time it is crossed. We denote the probability distribution of c_{ij} by p_{ij} . Throughout the paper, we also assume that the arc cost distributions have compact supports. This is a perfectly reasonable assumption in many practical settings, such as in transportation networks.

ASSUMPTION 3. $\forall (i, j) \in \mathcal{A}$, p_{ij} has compact support included in $[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]$ with $\delta_{ij}^{\inf} > 0$ and $\delta_{ij}^{\sup} < \infty$. Thus $\delta^{\inf} = \min_{(i,j) \in \mathcal{A}} \delta_{ij}^{\inf} > 0$ and $\delta^{\sup} = \max_{(i,j) \in \mathcal{A}} \delta_{ij}^{\sup} < \infty$.

Assumption 3 is motivated by computational considerations, see Section 3.2.2, but is also substantially needed when proving theoretical properties satisfied by optimal solutions to (1), and when analyzing the complexity of the proposed algorithms.

2.2. Distributionally robust problem

One of the major limitations of the approach described in Section 2.1 is that it requires a full description of the uncertainty. Under Assumptions 1 and 2, this is equivalent to having access to the exact arc cost probability distributions. Yet, in practice, we often only have access to a limited number of realizations of the random variables c_{ij} . In these circumstances, it is tempting to estimate empirical arc cost distributions and to take them as input to problem (1). However, estimating accurately a distribution with samples drawn from realizations usually requires a very large sample size, and our experimental evidence suggests that, as a result, the corresponding solutions may perform poorly when only few samples are available, as we will see in Section 5. To mitigate the impact of the lack of information on the arc cost distributions, we adopt a distributionally robust point of view where, for each arc $(i, j) \in \mathcal{A}$, we assume that p_{ij} is only known to lie in an ambiguity set \mathcal{P}_{ij} . We make the following assumption on these ambiguity sets throughout the paper.

ASSUMPTION 4. $\forall (i, j) \in \mathcal{A}$, \mathcal{P}_{ij} is not empty, closed for the weak topology, and a subset of $\mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}])$, the set of probability measures on $[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]$.

The last part of Assumption 4 is a natural extension of Assumption 3, and is essential for computational tractability, see Section 4. The robust counterpart of (1) for an ambiguity-averse user is then given by:

$$\sup_{\pi \in \Pi} \inf_{\forall (i, j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}}[f(T - X_{\pi})], \quad (2)$$

where the notation \mathbf{p} refers to the fact that the costs $(c_{ij})_{(i, j) \in \mathcal{A}}$ are independent and distributed according to $(p_{ij})_{(i, j) \in \mathcal{A}}$.

As a byproduct of the results obtained for the nominal problem in Section 3.1, (2) can be equivalently viewed as a distributionally robust MDP in the extended space state $(i, \tau) \in \mathcal{V} \times \mathbb{R}_+$ where i is the current location and τ is the total cost spent so far and where the transition probabilities from any state (i, τ) to any state (j, τ') , for $j \in \mathcal{V}(i)$ and $\tau' \geq \tau$, are only known to jointly lie in a global ambiguity set. As shown in Wiesemann et al. (2013), the tractability of a distributionally robust MDP hinges on the decomposability of the global ambiguity set as a Cartesian product over the space state of individual ambiguity sets, a property coined as *rectangularity*. While the global ambiguity set of (2) is rectangular with respect to our original state space \mathcal{V} , it is not with respect to the extended space space $\mathcal{V} \times \mathbb{R}_+$. Thus, we are led to enlarge our ambiguity set to make it rectangular and consider a robust relaxation of (2). This boils down to allowing the arc cost distributions to vary in their respective ambiguity sets as a function of the total cost spent so far. This approach leads to the following choice for our robust formulation associated with an ambiguity-averse user:

$$\sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i, j) \in \mathcal{A}, p_{ij}^{\tau} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^{\tau}}[f(T - X_{\pi})], \quad (3)$$

where the notation \mathbf{p}^τ refers to the fact that, for any arc $(i, j) \in \mathcal{A}$, the costs $(c_{ij}^\tau)_{\tau \geq 0}$ are independent and distributed according to $(p_{ij}^\tau)_{\tau \geq 0}$. Note that when Assumption 2 is relaxed, we have a different ambiguity set for each pair $((i, j), \tau) \in \mathcal{A} \times \mathbb{R}_+$, which is denoted in this case by \mathcal{P}_{ij}^τ , and (3) is precisely the robust counterpart of (1) as opposed to a robust relaxation, see Section 6.2. Also observe that (3) reduces to (1) when the ambiguity sets are singleton, i.e. $\mathcal{P}_{ij} = \{p_{ij}\}$. In the sequel, we focus on (3) and refer to this optimization problem as the robust problem. But we will also investigate the performance of an optimal solution to (3) with respect to the optimization problem (2), both from a theoretical standpoint in Section 4.3.2, and from a practical standpoint in Section 5. Finally note that we consider general ambiguity sets satisfying Assumption 4 when we study the theoretical properties of (3). However, for tractability purposes, the solution procedure that we develop in Section 4.3.3 only applies to ambiguity sets defined by confidence intervals on piecewise affine statistics, such as the mean, the absolute mean deviation, or any quantile. We refer to Section 4.3.2 for a discussion on the modeling power of these ambiguity sets and on how to build them with samples. Similarly as for the nominal problem, we will also restrict our attention to risk functions satisfying natural properties meant to prevent infinite cycling in Theorem 2 of Section 4.1.

3. Theoretical and computational analysis of the nominal problem

3.1. Characterization of optimal policies

Perhaps the most important property of (1) is that *Bellman's Principle of Optimality* can be shown to hold. Specifically, for any history of the process $h \in \mathcal{H}$, an optimal strategy to (1) must also be an optimal strategy to the subproblem of minimizing the risk function given this history. Otherwise, we could modify this strategy for this particular history and take it to be an optimal strategy for this subproblem. This operation could only increase the objective function of the optimization problem (1), which would contradict the optimality of the strategy.

Another, less obvious, interesting feature of (1) is that, even for perfectly natural risk functions $f(\cdot)$, making decisions according to an optimal strategy may lead to cycle back to a previously visited location. This may happen, for instance, when the objective is to maximize the probability of completion within budget, see Samaranayake et al. (2012b), and their example can be adapted when the objective is to minimize the expected budget overrun, see Figure 1. While counter-intuitive at first, the existence of loops is a direct consequence of the stochasticity of the costs when the decision maker is concerned about the risk of going over budget, as illustrated in Figure 1. On the other hand, the existence of infinitely many loops is particularly troublesome from a modeling perspective as it would imply that a user traveling through \mathcal{V} following the optimal strategy may get at a location $i \neq d$ having already spent an arbitrarily large budget with positive probability. Furthermore, infinite cycling is also problematic from a computational standpoint because describing an optimal strategy




Figure 1 Existence of loops. If the initial budget is $T = 8$ and the risk function is $f(t) = t \cdot 1_{t \leq 0}$, the optimal strategy to travel from s to d is to go to a first. This is because going to d directly incurs an expected delay of 0.1, while going to a first and then planning to go to d incurs an expected delay of 0.01. If we end up getting a cost $c_{sa} = 5$ on the way to a , then, performing a similar analysis, the optimal strategy is to go back to s .

would require unlimited storage capacity. We argue that infinite cycling arises only when the risk function is poorly chosen. This is obvious when $f(t) = -t \cdot 1_{t \leq 0}$, which corresponds to maximizing the expected budget overrun, but we stress that it is not merely a matter of monotonicity. Infinite cycling may occur even if $f(\cdot)$ is increasing as we highlight in Example 1.

EXAMPLE 1. Consider the simple directed graph of Figure 2a and the risk function $f(\cdot)$ illustrated in Figure 2b. $f(\cdot)$ is defined piecewise, alternating between concavity and convexity on intervals of size T^* and the same pattern is repeated every $2T^*$. This means that, for this particular objective, the attitude towards risk keeps fluctuating as the budget decreases, from being risk-averse when $f(\cdot)$ is locally concave to being risk-seeking when $f(\cdot)$ is locally convex. Now take $\delta^{\text{inf}} \ll 1$, $\epsilon \ll 1$ and $T^* > 3$ and consider finding a strategy to get to d starting from s with initial budget T which we choose to take at a point where $f(\cdot)$ switches from being concave to being convex, see Figure 2b. Going straight to d incurs an expected objective value of $f(T - 2) < \frac{1}{2}f(T - 1) + \frac{1}{2}f(T - 3)$ and we can make this gap arbitrarily large by properly defining $f(\cdot)$. Therefore, by taking ϵ and δ^{inf} small enough, going to a first is optimal. With probability $\epsilon > 0$, we arrive at a with a remaining budget of $T - T^*$. Afterwards, the situation is reversed as we are willing to take as little risk as possible and the corresponding optimal solution is to go back to s . With probability ϵ , we arrive at s with a budget of $T - 2T^*$ and we are back in the initial situation, showing the existence of infinite cycling.

In light of Example 1, we identify a set of sufficient asymptotic conditions on $f(\cdot)$ ruling out the possibility of infinite cycling.

THEOREM 1. *Case 1: If there exists T_1 such that either:*

- (a) *$f(\cdot)$ is increasing, concave, and C^2 on $(-\infty, T_1)$ and such that $\frac{f''}{f'} \rightarrow_{-\infty} 0$,*




Figure 2 Existence of infinite cycling from Example 1.

(b) $f(\cdot)$ is C^1 on $(-\infty, T_1)$ and $\lim_{-\infty} f'$ exists, is positive, and is finite, then there exists T_f such that, for any $T \geq 0$ and as soon as the total cost spent so far is larger than $T - T_f$, any optimal policy to (1) follows the shortest-path tree rooted at d with respect to the mean arc costs, which we denote by \mathcal{T} .

Case 2: If there exists T_f such that the support of $f(\cdot)$ is included in $[T_f, \infty)$, then following \mathcal{T} is optimal as soon as the total cost spent so far is larger than $T - T_f$.

For a node i , $\mathcal{T}(i)$ refers to the set of immediate successors of i in \mathcal{T} . The proof is deferred to the online supplement, Section B.1.

Observe that, in addition to not being concave, the choice of $f(\cdot)$ in Example 1 does not satisfy property (b) as $f'(\cdot)$ is $2T^*$ -periodic. An immediate consequence of Theorem 1 is that an optimal strategy to (1) does not include any loop as soon as the total cost spent so far is larger than $T - T_f$. Since each arc has a positive minimum cost, this rules out infinite cycling. The parameter T_f can be computed through direct reasoning on the risk function $f(\cdot)$ or by inspecting the proof of Theorem 1. Remark that any polynomial of even degree with a negative leading coefficient satisfies condition (a) of Theorem 1. Examples of valid objectives include maximization of the probability of completion within budget $f(t) = 1_{t \geq 0}$ with $T_f = 0$, minimization of the budget overrun $f(t) = t \cdot 1_{t \leq 0}$ with $T_f = 0$, and minimization of the squared budget overrun $f(t) = -t^2 \cdot 1_{t \leq 0}$ with

$$T_f = -\frac{|\mathcal{V}| \cdot \delta^{\sup} \cdot \max_{i \in \mathcal{V}} M_i}{2 \cdot \min_{i \neq d} \min_{j \in \mathcal{V}(i), j \notin \mathcal{T}(i)} \{\mathbb{E}[c_{ij}] + M_j - M_i\}},$$

where M_i is the minimum expected cost to go from i to d and with the convention that the minimum of an empty set is equal to ∞ . When $f(\cdot)$ is increasing but does not satisfy condition (a) or (b), the optimal strategy may follow a different shortest-path tree. For instance, if $f(t) = -\exp(-t)$, the optimal policy is to follow the shortest path to d with respect to $(\log(\mathbb{E}[\exp(c_{ij})]))_{(i,j) \in \mathcal{A}}$. Conversely, if $f(t) = \exp(t)$, the optimal policy is to follow the shortest path to d with respect to

$(-\log(\mathbb{E}[\exp(-c_{ij})]))_{(i,j) \in \mathcal{A}}$. For these reasons, proving that an optimal strategy to (1) does not include infinitely many loops when $f(\cdot)$ does not satisfy the assumptions of Theorem 1 requires objective-specific (and possibly graph-specific) arguments. To illustrate this last point, observe that the conclusion of Theorem 1 always holds for a graph consisted of a single simple path regardless of the definition of $f(\cdot)$, even if this function is decreasing. Hence, the assumptions of Theorem 1 are not necessary in general to prevent infinite cycling but restricting our attention to this class of risk functions enables us to study the problem in a generic fashion and to develop a general-purpose algorithm in Section 3.2.

Another remarkable property of (1) is that it can be equivalently formulated as a MDP in the extended space state $(i, t) \in \mathcal{V} \times (-\infty, T]$ where i is the current location and t is the remaining budget. As a result, standard techniques for MDPs can be applied to show that there exists an optimal *Markov* policy π_f^* which is a mapping from the current location and the remaining budget to the next node to visit. Furthermore, the optimal *Markov* policies are characterized by the dynamic programming equation:

$$\begin{aligned} u_d(t) &= f(t) & t \leq T \\ u_i(t) &= \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega) d\omega & i \neq d, t \leq T \\ \pi_f^*(i, t) &\in \arg \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega) d\omega & i \neq d, t \leq T, \end{aligned} \quad (4)$$

where $\mathcal{V}(i) = \{j \in \mathcal{V} \mid (i, j) \in \mathcal{A}\}$ refers to the set of immediate successors of i in \mathcal{G} and $u_i(t)$ is the expected objective-to-go when leaving $i \in \mathcal{V}$ with remaining budget t . The interpretation of (4) is simple. At each node $i \in \mathcal{V}$, and for each potential remaining budget t , the decision maker should pick the outgoing edge (i, j) that yields the maximum expected objective-to-go if acting optimally thereafter.

PROPOSITION 1. *Under the same assumptions as in Theorem 1, any Markov policy solution to (4) is an optimal strategy for (1).*

The proof is deferred to the online supplement, Section B.2.

3.2. Solution methodology

In order to solve (1), we use Proposition 1 and compute a *Markov* policy solution to the dynamic program (4). We face two main challenges when we carry out this task. First, (4) is a continuous dynamic program. To solve this program numerically, we approximate the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ by piecewise constant functions, as detailed in Section 3.2.1. Second, as illustrated in Figure 1 of Section 3.1, an optimal *Markov* strategy solution to (4) may contain loops. Hence, in the presence of a cycle in \mathcal{G} , say $i \rightarrow j \rightarrow i$, observe that computing $u_i(t)$ requires to know the value of $u_j(t)$ which in turns

depends on $u_i(t)$. As a result, it is a-priori unclear how to solve (4) without resorting to value or policy iteration. We explain how to sidestep this difficulty and construct efficient label-setting algorithms in Section 3.2.2. In particular, using these algorithms, we can compute:

- an optimal solution to (1) in $O(|\mathcal{A}| \cdot \frac{T-T_f}{\Delta t} \cdot \log^2(\frac{\delta^{\text{sup}}}{\Delta t}) + |\mathcal{V}|^2 \cdot \frac{\delta^{\text{sup}}}{\Delta t} \cdot \log(|\mathcal{V}| \cdot \frac{\delta^{\text{sup}}}{\Delta t}))$ computation time when the arc costs only take on values that are multiple of $\Delta t > 0$ and for any risk function $f(\cdot)$ satisfying Theorem 1. This simplifies to $O(|\mathcal{A}| \cdot \frac{T}{\Delta t} \cdot \log^2(\frac{\delta^{\text{sup}}}{\Delta t}))$ when the objective is to maximize the probability of completion within budget,

- an ϵ -approximate solution to (1) in

$$O\left(\frac{(|\mathcal{V}| + \frac{T-T_f}{\delta^{\text{inf}}})^2}{\epsilon} \cdot \left[|\mathcal{A}| \cdot (T-T_f) \cdot \log^2\left(\frac{(|\mathcal{V}| + \frac{T-T_f}{\delta^{\text{inf}}}) \cdot \delta^{\text{sup}}}{\epsilon}\right) + |\mathcal{V}|^2 \cdot \delta^{\text{sup}} \cdot \log\left(\frac{(|\mathcal{V}| + \frac{T-T_f}{\delta^{\text{inf}}}) \cdot |\mathcal{V}| \cdot \delta^{\text{sup}}}{\epsilon}\right) \right] \right)$$

computation time when the risk function is Lipschitz on compact sets.

3.2.1. Discretization scheme For each node $i \in \mathcal{V}$, we approximate $u_i(\cdot)$ by a piecewise constant function $u_i^{\Delta t}(\cdot)$ of uniform stepsize Δt . Under the conditions of Theorem 1, we only need to approximate $u_i(\cdot)$ for a remaining budget larger than $k_i^{\min} \cdot \Delta t$, for $k_i^{\min} = \left\lfloor \frac{T_f - (|\mathcal{V}| - \text{level}(i, \mathcal{T}) + 1) \cdot \delta^{\text{sup}}}{\Delta t} \right\rfloor$, where $\text{level}(i, \mathcal{T})$ is defined as the level of node i in the rooted tree \mathcal{T} , i.e. the number of parent nodes of i in \mathcal{T} plus one. This is because, following the shortest path tree \mathcal{T} once the remaining budget drops below T_f , we can never get to state i with remaining budget less than $k_i^{\min} \cdot \Delta t$. We use the approximation:

$$\begin{aligned} u_i^{\Delta t}(t) &= u_i^{\Delta t}\left(\left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t\right) & i \in \mathcal{V}, t \in [k_i^{\min} \cdot \Delta t, T] \\ \pi^{\Delta t}(i, t) &= \pi^{\Delta t}\left(i, \left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t\right) & i \neq d, t \in [k_i^{\min} \cdot \Delta t, T], \end{aligned} \quad (5)$$

and the values at the mesh points are determined by the set of equalities:

$$\begin{aligned} u_d^{\Delta t}(k \cdot \Delta t) &= f(k \cdot \Delta t) & k = k_d^{\min}, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ u_i^{\Delta t}(k \cdot \Delta t) &= \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k = \left\lfloor \frac{T_f}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ \pi^{\Delta t}(i, k \cdot \Delta t) &\in \arg \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k = \left\lfloor \frac{T_f}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ u_i^{\Delta t}(k \cdot \Delta t) &= \max_{j \in \mathcal{T}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k = k_i^{\min}, \dots, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor - 1 \\ \pi^{\Delta t}(i, k \cdot \Delta t) &\in \arg \max_{j \in \mathcal{T}(i)} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k = k_i^{\min}, \dots, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor - 1. \end{aligned} \quad (6)$$

Notice that for $t \leq T_f$, we rely on Theorem 1 and only consider, for each node $i \neq d$, the immediate neighbors of i in \mathcal{T} . This is of critical importance to be able to solve (6) with a label-setting algorithm, see Section 3.2.2. The next result provides insight into the quality of the policy $\pi^{\Delta t}$ as an approximate solution to (1).

PROPOSITION 2. Consider a solution to the global discretization scheme (5) and (6), $(\pi^{\Delta t}, (u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}})$. We have:

1. If $f(\cdot)$ is non-decreasing, the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge pointwise almost everywhere to $(u_i(\cdot))_{i \in \mathcal{V}}$ as $\Delta t \rightarrow 0$,
2. If $f(\cdot)$ is continuous, the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge uniformly to $(u_i(\cdot))_{i \in \mathcal{V}}$ and $\pi^{\Delta t}$ is a $o(1)$ -approximate optimal solution to (1) as $\Delta t \rightarrow 0$,
3. If $f(\cdot)$ is Lipschitz on compact sets (e.g. if $f(\cdot)$ is C^1), the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge uniformly to $(u_i(\cdot))_{i \in \mathcal{V}}$ at speed Δt and $\pi^{\Delta t}$ is a $O(\Delta t)$ -approximate optimal solution to (1) as $\Delta t \rightarrow 0$,
4. If $f(t) = 1_{t \geq 0}$ and the distributions $(p_{ij})_{(i,j) \in \mathcal{A}}$ are continuous, the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge uniformly to $(u_i(\cdot))_{i \in \mathcal{V}}$ and $\pi^{\Delta t}$ is a $o(1)$ -approximate optimal solution to (1) as $\Delta t \rightarrow 0$.

The proof is deferred to the online supplement, Section B.3.

If the distributions $(p_{ij})_{(i,j) \in \mathcal{A}}$ are discrete and $f(\cdot)$ is piecewise constant, an exact optimal solution to (1) can be computed by appropriately choosing a different discretization length for each node. In this paper, we focus on discretization schemes with a uniform stepsize Δt for mathematical convenience. We stress that choosing adaptively the discretization length can improve the quality of the approximation for the same number of computations, see Hoy and Nikolova (2015).

3.2.2. Solution procedures The key observation enabling the development of label-setting algorithms to solve (4) is made by Samaranayake et al. (2012b). They note that, when the risk function is the probability of completion within budget, $u_i(t)$ can be computed for $i \in \mathcal{V}$ and $t \leq T$ as soon as the values taken by $u_j(\cdot)$ on $(-\infty, t - \delta^{\text{inf}}]$ are available for all neighboring nodes $j \in \mathcal{V}(i)$ since $p_{ij}(\omega) = 0$ for $\omega \leq \delta^{\text{inf}}$ under Assumption 3. They propose a label-setting algorithm which consists in computing the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ block by block, by interval increments of size δ^{inf} . After the following straightforward initialization step: $u_i(t) = 0$ for $t \leq 0$ and $i \in \mathcal{V}$, they first compute $(u_i(\cdot)_{[0, \delta^{\text{inf}}]})_{i \in \mathcal{V}}$, then $(u_i(\cdot)_{[0, 2 \cdot \delta^{\text{inf}}]})_{i \in \mathcal{V}}$ and so on to eventually derive $(u_i(\cdot)_{[0, T]})_{i \in \mathcal{V}}$. While this incremental procedure can still be applied for general risk functions, the initialization step gets tricky if $f(\cdot)$ does not have a one-sided compact support of the type $[a, \infty)$. Theorem 1 is crucial in this respect because the shortest-path tree \mathcal{T} induces an ordering of the nodes to initialize the collection of functions $(u_i(\cdot))_{i \in \mathcal{V}}$ for remaining budgets smaller than T_f . The functions can subsequently be computed for larger budgets using the incremental procedure outlined above. To be specific, we solve (6) in three steps. First, we compute T_f (defined in Theorem 1). Inspecting the proof of Theorem 1, observe that T_f only depends on few parameters, namely the risk function $f(\cdot)$, the expected arc costs, and the maximum arc costs. Next, we compute the values $u_i^{\Delta t}(k \cdot \Delta t)$ for $k \in \{k_i^{\min}, \dots, \lfloor \frac{T_f}{\Delta t} \rfloor - 1\}$ starting at node $i = d$ and traversing the tree \mathcal{T} in a breadth-first fashion using fast Fourier

transforms with complexity $O(|\mathcal{V}|^2 \cdot \frac{\delta^{\sup}}{\Delta t} \cdot \log(|\mathcal{V}| \cdot \frac{\delta^{\sup}}{\Delta t}))$. Note that this step can be made to run significantly faster for specific risk functions, e.g. for the probability of completion within budget where $u_i^{\Delta t}(k \cdot \Delta t) = 0$ for $k < \lfloor \frac{T_f}{\Delta t} \rfloor$ and any $i \in \mathcal{V}$. Finally, we compute the values $u_i^{\Delta t}(k \cdot \Delta t)$ for $k \in \{\lfloor \frac{T_f}{\Delta t} \rfloor + m \cdot \lfloor \frac{\delta^{\inf}}{\Delta t} \rfloor, \dots, \lfloor \frac{T_f}{\Delta t} \rfloor + (m+1) \cdot \lfloor \frac{\delta^{\inf}}{\Delta t} \rfloor\}$ for all nodes $i \in \mathcal{V}$ by induction on m .

Complexity analysis. The description of the last step of the label-setting approach leaves out one detail that has a dramatic impact on the runtime complexity. We need to specify how to compute the convolution products arising in (6) for $k \geq \lfloor \frac{T_f}{\Delta t} \rfloor$, keeping in mind that, for any node $i \in \mathcal{V}$, the values $u_i^{\Delta t}(k \cdot \Delta t)$ for $k \in \{\lfloor \frac{T_f}{\Delta t} \rfloor, \dots, \lfloor \frac{T}{\Delta t} \rfloor\}$ become available online by chunks of length $\lfloor \frac{\delta^{\inf}}{\Delta t} \rfloor$ as the label-setting algorithm progresses. A naive implementation consisting in applying the pointwise definition of convolution products has a runtime complexity $O(|\mathcal{A}| \cdot \frac{(T-T_f) \cdot (\delta^{\sup} - \delta^{\inf})}{(\Delta t)^2})$. Using fast Fourier transforms for each chunk brings down the complexity to $O(|\mathcal{A}| \cdot \frac{(T-T_f)}{\Delta t} \cdot \frac{\delta^{\sup}}{\delta^{\inf}} \cdot \log(\frac{\delta^{\sup}}{\Delta t}))$. Applying another online scheme developed in Dean (2010) and Samaranayake et al. (2012a), based on the idea of zero-delay convolution, leads to a worst-case complexity $O(|\mathcal{A}| \cdot \frac{(T-T_f)}{\Delta t} \cdot \log^2(\frac{\delta^{\sup}}{\Delta t}))$. Numerical evidence suggest that this last implementation significantly speeds up the computations, see Samaranayake et al. (2012a).

4. Theoretical and computational analysis of the robust problem

4.1. Characterization of optimal policies

The properties satisfied by optimal solutions to the nominal problem naturally extend to their robust counterparts, which we recall are defined as optimal solutions to (3). In fact, all the results derived in this section are strict generalizations of those obtained in Section 3.1 for singleton ambiguity sets. We point out that the rectangularity of the global ambiguity set is essential for the results to carry over to the robust setting as it guarantees that *Bellman's Principle of Optimality* continue to hold, which is an absolute prerequisite for computational tractability.

Similarly as what we have seen for the nominal problem, infinite cycling might occur in the robust setting, depending on the risk function at hand. This difficulty can be shown not to arise under the same conditions on $f(\cdot)$ as for the nominal problem.

THEOREM 2. *Case 1: If there exists T_1 such that either:*

- (a) *$f(\cdot)$ is increasing, concave, and C^2 on $(-\infty, T_1)$ and such that $\frac{f''}{f'} \rightarrow_{-\infty} 0$,*
- (b) *$f(\cdot)$ is C^1 on $(-\infty, T_1)$ and $\lim_{-\infty} f'$ exists, is positive, and is finite,*

then there exists T_f^r such that, for any $T \geq 0$ and as soon as the total cost spent so far is larger than $T - T_f^r$, any optimal policy solution to (3) follows the shortest-path tree rooted at d with respect to the worst-case mean arc costs, i.e. $(\max_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[X])_{(i,j) \in \mathcal{A}}$, which we denote by \mathcal{T}^r .

Case 2: If there exists T_f such that the support of $f(\cdot)$ is included in $[T_f, \infty)$, then following \mathcal{T}^r is optimal as soon as the total cost spent so far is larger than $T - T_f^r$.

For a node i , $\mathcal{T}^r(i)$ refers to the set of immediate successors of node i in \mathcal{T}^r . The proof is deferred to the online supplement, Section B.4.

Interestingly, T_f^r is determined by the exact same procedure as T_f provided the expected arc costs are substituted with the worst-case expected costs. For instance, when $f(t) = -t^2 \cdot 1_{t \leq 0}$, we may take:

$$T_f^r = -\frac{|\mathcal{V}| \cdot \delta^{\sup} \cdot \max_{i \in \mathcal{V}} M_i}{2 \cdot \min_{i \neq d} \min_{j \in \mathcal{V}(i), j \notin \mathcal{T}^r(i)} \{\max_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[X] + M_j - M_i\}},$$

where M_i is the worst-case minimum expected cost to go from i to d .

Last but not least, problem (3) can be formulated as a distributionally robust MDP in the extended space state $(i, t) \in \mathcal{V} \times (-\infty, T]$. As a result, one can show that there exists an optimal *Markov* policy $\pi_{f, \mathcal{P}}^*$ characterized by the dynamic programming equation:

$$\begin{aligned} u_d(t) &= f(t) & t \leq T \\ u_i(t) &= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega) d\omega & i \neq d, t \leq T \\ \pi_{f, \mathcal{P}}^*(i, t) &\in \arg \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega) d\omega & i \neq d, t \leq T, \end{aligned} \quad (7)$$

where $u_i(t)$ is the worst-case expected objective-to-go when leaving $i \in \mathcal{V}$ with remaining budget t . Observe that (7) only differs from (4) through the presence of the infimum over \mathcal{P}_{ij} .

PROPOSITION 3. *Any Markov policy solution to (7) is an optimal strategy for (3).*

The proof is deferred to the online supplement, Section B.5.

4.2. Tightness of the robust problem

The optimization problem (3) is a robust relaxation of (2) in the sense that, for any strategy $\pi \in \Pi$, we have:

$$\inf_{\forall (i, j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}}[f(T - X_\pi)] \geq \inf_{\forall \tau, \forall (i, j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^\tau}[f(T - X_\pi)].$$

We say that (2) and (3) are equivalent if they share the same optimal value and if there exists a common optimal strategy. For general risk functions, ambiguity sets, and graphs, (2) and (3) are not equivalent. In this section, we highlight several situations of interest for which (2) and (3) happen to be equivalent and we bound the gap between the optimal values of (2) and (3) for a subclass of risk functions. In this paper, we solve (3) instead of (2) for computational tractability, irrespective of whether or not (2) and (3) are equivalent. Hence, the results presented in this section are included mainly for illustrative purposes, i.e. we do not impose further restrictions on the risk function or the ambiguity sets here.

Equivalence of (2) and (3). As a simple first example, observe that when $f(\cdot)$ is non-decreasing and $\mathcal{P}_{ij} = \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}])$, both (2) and (3) reduce to a standard robust approach where the goal is to find a path minimizing the sum of the worst-case arc costs. The following result identifies conditions of broader applicability when the decision maker is risk-seeking.

LEMMA 1. *Suppose that $f(\cdot)$ is convex and satisfies property (b) in Case 1 of Theorem 2 and that, for any arc $(i, j) \in \mathcal{V}$, the Dirac distribution supported at $\max_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[X]$ belongs to \mathcal{P}_{ij} . Then, (2) and (3) are equivalent.*

The proof is deferred to the online supplement, Section B.7.

To illustrate Lemma 1, observe that the assumptions are satisfied for $f(t) = \exp(a \cdot t) + b \cdot t$, with a and b taken as positive values, and when the ambiguity sets are defined through confidence intervals on the expected costs, i.e. for any arc $(i, j) \in \mathcal{A}$:

$$\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \mathbb{E}_{X \sim p}[X] \in [\alpha_{ij}, \beta_{ij}]\},$$

with $\alpha_{ij} \leq \beta_{ij}$. Further note that adding upper bounds on the mean deviation or on higher order moments in the definition of the ambiguity sets does not alter the conclusion of Lemma 1. We move on to another situation of interest where (2) and (3) can be shown to be equivalent.

LEMMA 2. *Take $K \in \mathbb{N}$. Suppose that:*

- \mathcal{G} is a single-path graph,
- $f(\cdot)$ is C^{K+1} and $f^{(K+1)}(t) > 0 \forall t$ or $f^{(K+1)}(t) < 0 \forall t$,
- For any arc $(i, j) \in \mathcal{A}$:

$$\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \mathbb{E}_{X \sim p}(X) = m_{ij}^1, \dots, \mathbb{E}_{X \sim p}(X^K) = m_{ij}^K\},$$

where $m_{ij}^1, \dots, m_{ij}^K$ are non-negative.

Then (2) and (3) are equivalent.

The proof is deferred to the online supplement, Section B.8.

When \mathcal{G} is a single-path graph, the optimal value of (2) corresponds to the worst-case risk function when following this path, given that the arc cost distributions are only known to lie in the ambiguity sets. While it is a priori unclear how to compute this quantity, Proposition 4 of Section 4.3.1 establishes that the optimal value of (3) can be determined with arbitrary precision provided the inner optimization problems appearing in the discretization scheme of Section 4.3.1 can be computed numerically. Hence, even in this seemingly simplistic situation, the equivalence between (2) and (3) is an important fact to know as it has significant computational implications. Lemma 2 shows that, when the risk function is $(K+1)$ th order convex or concave and when the arc cost distributions are only known through the first K -order moments, (2) and (3) are in fact equivalent. For this particular class of ambiguity sets, the inner optimization problems of the discretization scheme of Section 4.3.1 can be solved using semidefinite programming, see Bertsimas and Popescu (2005).

Bounding the gap between the optimal values of (2) and (3). It turns out that, for a particular subclass of risk functions, we can bound the gap between the optimal values of (2) and (3) uniformly over all graphs and ambiguity sets.

LEMMA 3. Denote the optimal value of (2) (resp. (3)) by v^* (resp. v).

If there exists $\gamma, a > 0$ and β, b such that one of the following conditions holds:

- $\gamma \cdot t + \beta \geq f(t) \geq a \cdot t + b \quad \forall t \leq T,$
- $\gamma \cdot \exp(t) + \beta \geq f(t) \geq a \cdot \exp(t) + b \quad \forall t \leq T,$
- $-\gamma \cdot \exp(-t) + \beta \geq f(t) \geq -a \cdot \exp(-t) + b \quad \forall t \leq T,$

then $v^* \geq v \geq \frac{a}{\gamma} \cdot (v^* - \beta) + b$.

The proof is deferred to the online supplement, Section B.9.

4.3. Solution methodology

We proceed as in Section 3.2 and compute an approximate *Markov* policy solution to (7). The computational challenges faced when solving the nominal problem carry over to the robust counterpart, but with additional difficulties to overcome. Specifically, the continuity of the problem leads us to build a discrete approximation in Section 4.3.1 similar to the one developed for the nominal approach. We also extend the label-setting algorithm of Section 3.2.2 to tackle the potential existence of cycles at the beginning of Section 4.3.3. However, the presence of an inner optimization problem in (7) is a distinctive feature of the robust problem which poses a new computational challenge. As a result, and in contrast with the situation for the nominal problem where this optimization problem reduces to a convolution product, it is not a priori obvious how to solve the discretization scheme numerically, let alone efficiently. As can be expected, the exact form taken by the ambiguity sets has a major impact on the computational complexity of the inner optimization problem. In an effort to mitigate the computational burden, we restrict our attention to a subclass of ambiguity sets defined by confidence intervals on piecewise affine statistics in Section 4.3.2. While this simplification might seem restrictive, we show that this subclass displays significant modeling power. Finally, we develop two general-purpose algorithms in Section 4.3.3 for this particular subclass of ambiguity sets. The computational attractiveness of these approaches hinges on the existence of a data structure, presented in Section 4.3.4, maintaining the convex hull of a dynamic set of points efficiently. The mechanism behind this data structure can be regarded as the counterpart of the online fast Fourier scheme for the nominal approach. In particular, using the algorithms developed in this section, we can compute:

- an ϵ -approximate solution to (3) in

$$O\left(\frac{|\mathcal{A}| \cdot (T - T_f^r) + |\mathcal{V}|^2 \cdot \delta^{\sup}}{\Delta t} \cdot \log\left(\frac{\delta^{\sup} - \delta^{\inf}}{\Delta t}\right) \cdot \log\left(\frac{|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\inf}}}{\epsilon}\right)\right)$$

computation time when the arc costs only take on values that are multiple of $\Delta t > 0$ and for any continuous risk function $f(\cdot)$ satisfying Theorem 2. This also applies when the objective is to maximize the probability of completion within budget and even simplifies to $O(|\mathcal{A}| \cdot \frac{T}{\Delta t} \cdot \log(\frac{\delta^{\sup} - \delta^{\inf}}{\Delta t}) \cdot \log(\frac{T}{\epsilon \cdot \delta^{\inf}}))$,

- an ϵ -approximate solution to (3) in

$$O\left(\frac{(|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\inf}})^2 \cdot (|\mathcal{A}| \cdot (T - T_f^r) + |\mathcal{V}|^2 \cdot \delta^{\sup})}{\epsilon} \cdot \log\left(\frac{(|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\inf}}) \cdot (\delta^{\sup} - \delta^{\inf})}{\epsilon}\right) \cdot \log\left(\frac{|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\inf}}}{\epsilon}\right)\right)$$

computation time when the risk function is Lipschitz on compact sets.

4.3.1. Discretization scheme For each node $i \in \mathcal{V}$, we approximate $u_i(\cdot)$ by a piecewise affine continuous function $u_i^{\Delta t}(\cdot)$ of uniform stepsize Δt . This is in contrast with Section 3.2.1 where we use a piecewise constant approximation. This change is motivated by computational considerations. Essentially, the continuity of $u_i^{\Delta t}(\cdot)$ guarantees strong duality for the inner optimization problem appearing in (7). Similarly as for the nominal problem, we only need to approximate $u_i(\cdot)$ for a remaining budget larger than $k_i^{r,\min} \cdot \Delta t$, for $k_i^{r,\min} = \left\lfloor \frac{T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\sup}}{\Delta t} \right\rfloor$, where $\text{level}(i, \mathcal{T}^r)$ is the level of node i in \mathcal{T}^r . Specifically, we use the approximation:

$$\begin{aligned} u_i^{\Delta t}(t) &= (1 - \frac{t}{\Delta t} + \left\lfloor \frac{t}{\Delta t} \right\rfloor) \cdot u_i^{\Delta t}\left(\left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t\right) + \left(\frac{t}{\Delta t} - \left\lfloor \frac{t}{\Delta t} \right\rfloor\right) \cdot u_i^{\Delta t}\left(\left\lceil \frac{t}{\Delta t} \right\rceil \cdot \Delta t\right) \quad i \in \mathcal{V}, t \in [k_i^{r,\min} \cdot \Delta t, T] \\ \pi^{\Delta t}(i, t) &= \pi^{\Delta t}(i, \left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t) \quad i \neq d, t \in [k_i^{r,\min} \cdot \Delta t, T], \end{aligned} \quad (8)$$

and the values at the mesh points are determined by the set of equalities:

$$\begin{aligned} u_d^{\Delta t}(k \cdot \Delta t) &= f(k \cdot \Delta t) & k &= k_d^{r,\min}, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ u_i^{\Delta t}(k \cdot \Delta t) &= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k &= \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ \pi^{\Delta t}(i, k \cdot \Delta t) &\in \arg \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k &= \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor \\ u_i^{\Delta t}(k \cdot \Delta t) &= \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k &= k_i^{r,\min}, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor - 1 \\ \pi^{\Delta t}(i, k \cdot \Delta t) &\in \arg \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot u_j^{\Delta t}(k \cdot \Delta t - \omega) d\omega & i \neq d, k &= k_i^{r,\min}, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor - 1. \end{aligned} \quad (9)$$

As we did for the nominal problem, we can quantify the quality of $\pi^{\Delta t}$ as an approximate solution to (3) as a function of the regularity of the risk function.

PROPOSITION 4. *Consider a solution to the global discretization scheme (8) and (9), $(\pi^{\Delta t}, (u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}})$. We have:*

1. *If $f(\cdot)$ is non-decreasing, the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge pointwise almost everywhere to $(u_i(\cdot))_{i \in \mathcal{V}}$ as $\Delta t \rightarrow 0$.*

2. If $f(\cdot)$ is continuous, the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge uniformly to $(u_i(\cdot))_{i \in \mathcal{V}}$ and $\pi^{\Delta t}$ is a $o(1)$ -approximate optimal solution to (3) as $\Delta t \rightarrow 0$.
3. If $f(\cdot)$ is Lipschitz on compact sets (e.g. if $f(\cdot)$ is C^1), the functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ converge uniformly to $(u_i(\cdot))_{i \in \mathcal{V}}$ at speed Δt and $\pi^{\Delta t}$ is a $O(\Delta t)$ -approximate optimal solution to (3) as $\Delta t \rightarrow 0$.

The proof is deferred to the online supplement, Section B.6.

4.3.2. Ambiguity sets For computational tractability, we restrict our attention to the following subclass of ambiguity sets.

DEFINITION 1. For any arc $(i, j) \in \mathcal{A}$:

$$\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \mathbb{E}_{X \sim p}[g_q^{ij}(X)] \in [\alpha_q^{ij}, \beta_q^{ij}], q = 1, \dots, Q_{ij}\},$$

where:

- the functions $(g_q^{ij}(\cdot))_{q=1, \dots, Q_{ij}}$ are piecewise affine with a finite number of pieces on $[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]$ and such that \mathcal{P}_{ij} is closed for the weak topology,
- $-\infty \leq \alpha_q^{ij} \leq \beta_q^{ij} \leq \infty$ for $q = 1, \dots, Q_{ij}$.

Note that Definition 1 allows to model one-sided constraints by either taking $\alpha_q^{ij} = -\infty$ or $\beta_q^{ij} = \infty$. Moreover, we point out that the functions $(g_q^{ij}(\cdot))_{q=1, \dots, Q_{ij}}$ need not be continuous to guarantee closeness of \mathcal{P}_{ij} . For instance, the constraints $\mathbb{E}_{X \sim p}[1_{X \in S}] \leq \beta$ and $\mathbb{E}_{X \sim p}[1_{X \in S'}] \geq \beta$, for S (resp. S') an open (resp. a closed) set, are perfectly valid. In terms of modeling power, Definition 1 allows to have constraints on standard statistics, such as the mean value, the mean absolute deviation, and the median, but also to capture distributional asymmetry, through constraints on any quantile or of the type $\mathbb{E}_{X \sim p}[X \cdot 1_{X > \theta}] \leq \beta$, and to incorporate higher-order information, e.g. the variance or the skewness, since continuous functions can be approximated arbitrarily well by piecewise affine functions on a compact set. Finally, observe that Definition 1 also allows to model the situation where c_{ij} only takes values in a prescribed finite set S through the constraint $\mathbb{E}_{X \sim p}[1_{X \in S}] \geq 1$.

Data-driven ambiguity sets. Ambiguity sets of the form introduced in Definition 1 can be built using a combination of prior knowledge and historical data. To illustrate, suppose that, for any arc $(i, j) \in \mathcal{A}$, we have observed n_{ij} samples drawn from the corresponding arc cost distribution. Setting aside computational aspects, there is an inherent trade-off at play when designing ambiguity sets with this empirical data: using more statistics and/or narrowing the confidence intervals will improve the quality of the guarantee on the risk function provided by the robust approach, but will, on the other hand, deteriorate the probability that this guarantee holds. Assuming we are set on which statistics to use, the trade-off is simple to resolve as far as confidence intervals are concerned. Using

Hoeffding's and Boole's inequalities, the confidence interval for statistics q of arc (i, j) should be centered at the empirical average and have width ϵ_q^{ij} determined by:

$$\frac{\epsilon_q^{ij}}{\max_{[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]} g_q^{ij} - \min_{[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]} g_q^{ij}} = \sqrt{\frac{\log(\frac{2}{\epsilon} \cdot \sum_{(i,j) \in \mathcal{A}} Q_{ij})}{2n_{ij}}}$$

in order to achieve a probability $1 - \epsilon$ that the guarantee holds. Choosing which statistics to use is a more complex endeavor. It is not even clear whether using more statistics is beneficial since the confidence intervals jointly expand with $\sum_{(i,j) \in \mathcal{A}} Q_{ij}$. Numerical evidence presented in Section 5 suggests that low-order statistics, such as the mean, tend to be more informative when only few samples are available. Conversely, as sample sizes get very large, incorporating higher-order information seem to improve the quality of the strategy derived. In the limit where the statistics can be computed exactly, we should use as many statistics as possible. This observation is supported by the following lemma.

LEMMA 4. *For any arc $(i, j) \in \mathcal{A}$, consider $(\mathcal{P}_{ij}^k)_{k \in \mathbb{N}}$, a sequence of nested ambiguity sets satisfying Assumption 4. If $f(\cdot)$ is continuous, then the optimal value of the robust problem (3) when the uncertainty sets are taken as $(\mathcal{P}_{ij}^k)_{(i,j) \in \mathcal{A}}$ monotonically converges to the optimal value of (3) when the uncertainty sets are taken as $(\cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k)_{(i,j) \in \mathcal{A}}$ as $k \rightarrow \infty$.*

In particular, if $\cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k$ is a singleton for all arcs $(i, j) \in \mathcal{A}$, then the optimal value of the robust problem converges to the value of the nominal problem (1).

The proof is deferred to the online supplement, Section B.10.

Using the Weierstrass approximation theorem, observe that the second part of Lemma 4 applies in particular when the ambiguity sets \mathcal{P}_{ij}^k are defined by the first k -order moments.

4.3.3. Solution procedures We develop two general-purpose methods to compute a solution to the discretization scheme (9) for the class of ambiguity sets identified in Section 4.3.2. The first method, based on the ellipsoid algorithm, computes an ϵ -approximate solution to (9) with worst-case complexity:

$$O\left(\frac{|\mathcal{A}| \cdot (T - T_f^r) + |\mathcal{V}|^2 \cdot \delta^{\sup}}{\Delta t} \cdot \log\left(\frac{\delta^{\sup} - \delta^{\inf}}{\Delta t}\right) \cdot \log\left(\frac{|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\inf}}}{\epsilon}\right)\right),$$

provided $f(\cdot)$ is continuous and where the hidden factors are linear in the number of pieces of each statistic and polynomial in the number of statistics. We remind the reader that the complexity of solving the discretization scheme (6) for the nominal problem is $O(|\mathcal{A}| \cdot \frac{T - T_f}{\Delta t} \cdot \log^2(\frac{\delta^{\sup}}{\Delta t}) + |\mathcal{V}|^2 \cdot \frac{\delta^{\sup}}{\Delta t} \cdot \log(|\mathcal{V}| \cdot \frac{\delta^{\sup}}{\Delta t}))$ when using zero-delay convolution. While these bounds are not directly comparable because some of the parameters required to specify a robust instance are not relevant for a nominal instance and vice versa, we point out that they share many similarities, including the almost linear dependence on $\frac{1}{\Delta t}$. The second method, based on delayed column generation and warm starting

techniques, is more practical but has worst-case complexity exponential in $\frac{1}{\Delta t}$. We stress that none of these approaches can be used to solve the nominal problem as the latter is not a particular case of the robust problem for the restricted class of ambiguity sets defined in Section 4.3.2. Indeed, characterizing a single distribution generally requires infinitely many moment constraints.

Label-setting approach. To cope with the potential existence of cycles, we remark that the label-setting approach developed for the nominal approach trivially extends to the robust setting. Similarly as for the nominal problem, we proceed in three steps to solve (9). First, we compute T_f^r . Next, we compute the values $u_i^{\Delta t}(k \cdot \Delta t)$ for $k \in \{k_i^{r,\min}, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor - 1\}$ starting at node $i = d$ and traversing the tree \mathcal{T}^r in a breadth-first fashion. Finally, we compute the values $u_i^{\Delta t}(k \cdot \Delta t)$ for $k \in \left\{ \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor + m \cdot \left\lfloor \frac{\delta^{\inf}}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor + (m+1) \cdot \left\lfloor \frac{\delta^{\inf}}{\Delta t} \right\rfloor \right\}$ for all nodes $i \in \mathcal{V}$ by induction on m . Of course, an efficient procedure solving the inner optimization problem of (9) is a prerequisite for carrying out the last two steps. This will be our focus in the remainder of this section.

Solving the Inner Optimization Problem. Consider any arc $(i, j) \in \mathcal{A}$. We need to solve, at each step $k \in \{k_i^{r,\min}, \dots, \left\lfloor \frac{T}{\Delta t} \right\rfloor\}$, the optimization problem:

$$\begin{aligned} & \inf_{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}])} \mathbb{E}_{X \sim p}[u_j^{\Delta t}(k \cdot \Delta t - X)] \\ & \text{subject to} \quad \mathbb{E}_{X \sim p}[g_q^{ij}(X)] \in [\alpha_q^{ij}, \beta_q^{ij}] \quad q = 1, \dots, Q_{ij}. \end{aligned} \quad (10)$$

Since the set of non-negative measures on $[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]$ is a cone, (10) can be cast as a conic linear problem. As a result, standard conic duality theory applies and the optimal value of (10) can be equivalently computed by solving a dual optimization problem which turns out to be easier to study. For a thorough exposition of the duality theory of general conic linear problems, the reader is referred to Shapiro (2001). To simplify the presentation, we assume that $(\alpha_q^{ij})_{q=1, \dots, Q_{ij}}$ and $(\beta_q^{ij})_{q=1, \dots, Q_{ij}}$ are all finite quantities but this is by no means a limitation of our approach.

LEMMA 5. *The optimization problem (10) has the same optimal value as the semi-infinite linear program:*

$$\begin{aligned} & \sup_{\substack{z \in \mathbb{R} \\ y_1, \dots, y_{Q_{ij}} \in \mathbb{R} \\ x_1, \dots, x_{Q_{ij}} \in \mathbb{R}}} z + \sum_{q=1}^{Q_{ij}} (\alpha_q^{ij} \cdot x_q - \beta_q^{ij} \cdot y_q) \\ & \text{subject to} \quad z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot g_q^{ij}(\omega) \leq u_j^{\Delta t}(k \cdot \Delta t - \omega) \quad \forall \omega \in [\delta_{ij}^{\inf}, \delta_{ij}^{\sup}] \\ & \quad y_q, x_q \geq 0 \quad q = 1, \dots, Q_{ij}. \end{aligned} \quad (11)$$

Proof We take the Lagrangian dual of (10). Since $u_j^{\Delta t}(\cdot)$ is continuous by construction and \mathcal{P}_{ij} is not empty and compact by assumption, Proposition 3.1 in Shapiro (2001) shows that strong duality holds. \square

$$g_q^{ij}(\omega) = a_{q,r}^{ij} \cdot \omega + b_{q,r}^{ij} \quad \text{if } \omega \in I_r$$

LEMMA 6. *The semi-infinite linear program (11) can be reformulated as the following finite linear program:*

Proof Take $z, y_1, \dots, y_{Q_{ij}}, x_1, \dots, x_{Q_{ij}} \in \mathbb{R}$ and $r \in \{1, \dots, R_{ij}\}$. Since the function $\omega \rightarrow z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot g_q^{ij}(\omega)$ is affine on I_r , this function lies below the continuous piecewise affine function $u_j^{\Delta t}(k \cdot \Delta t - \cdot)$ on I_r if and only if it lies below $u_j^{\Delta t}(k \cdot \Delta t - \cdot)$ at every breakpoint of $u_j^{\Delta t}(k \cdot \Delta t - \cdot)$ on \bar{I}_r and at the boundary points of \bar{I}_r . Since the collection of intervals $(I_r)_{r=1, \dots, R_{ij}}$ forms a partition of $[\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]$, this establishes the claim. \square

$$z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot (a_{q,r}^{ij} \cdot l \cdot \Delta t + b_{q,r}^{ij}) \leq u_j^{\Delta t}((k-l) \cdot \Delta t)$$

does not limit the feasible region if $(l \cdot \Delta t, u_j^{\Delta t}((k-l) \cdot \Delta t))$ is not an extreme point of the upper convex hull of $\{(m \cdot \Delta t, u_j^{\Delta t}((k-m) \cdot \Delta t)), m = \lfloor \frac{\inf(I_r)}{\Delta t} \rfloor, \dots, \lceil \frac{\sup(I_r)}{\Delta t} \rceil\}$. Denote by $\mathcal{L}_{ij}^{k,r}$ the subset of integers l such that $(l \cdot \Delta t, u_j^{\Delta t}((k-l) \cdot \Delta t))$ is such an extreme point. Observe that the function

$$l \rightarrow u_j^{\Delta t}((k-l) \cdot \Delta t) - [z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot (a_{q,r}^{ij} \cdot l \cdot \Delta t + b_{q,r}^{ij})]$$

is convex on $\mathcal{L}_{ij}^{k,r}$, therefore a minimizer of this function can be found by binary search. As a result, all we need to be able to separate efficiently for the subset of constraints:

$$z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot (a_{q,r}^{ij} \cdot l \cdot \Delta t + b_{q,r}^{ij}) \leq u_j^{\Delta t}((k-l) \cdot \Delta t) \quad l = \left\lceil \frac{\inf(I_r)}{\Delta t} \right\rceil, \dots, \left\lfloor \frac{\sup(I_r)}{\Delta t} \right\rfloor$$

is a means to perform binary search on $\mathcal{L}_{ij}^{k,r}$ efficiently. We defer the presentation of a data structure designed for this purpose to Section 4.3.4 and make the following assumption to conclude the computational study.

ASSUMPTION 5. *For any two integers L, L' such that $\left\lceil \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rceil \leq L < L' \leq \left\lfloor \frac{\delta_{ij}^{\sup}}{\Delta t} \right\rfloor$, there exists a data structure that can maintain, dynamically as k increases from $k = k_i^{r,min}$ to $k = \lfloor \frac{T}{\Delta t} \rfloor$, a description of the upper convex hull of $\{(l \cdot \Delta t, u_j^{\Delta t}((k-l) \cdot \Delta t)), l = L, \dots, L'\}$ allowing to perform binary search on the first coordinate of the extreme points with a global complexity $O((\frac{T}{\Delta t} - k_i^{r,min}) \cdot \log(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}))$.*

Equipped with a data structure satisfying Assumption 5, the separation oracle has runtime complexity $O(\log(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}))$ given that there are at most $\left\lfloor \frac{\delta_{ij}^{\sup}}{\Delta t} \right\rfloor - \left\lceil \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rceil$ extreme points at any step k . Using the ellipsoid algorithm, we can compute the optimal value of (10) with precision ϵ in $O(\log(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}) \cdot \log(\frac{1}{\epsilon}))$ running time, where the hidden factors are polynomial in Q_{ij} and linear in R_{ij} . We point out that relying on a data structure satisfying Assumption 5 is critical to achieve this complexity: recomputing the upper convex hull from scratch at every time step k would increase the complexity to $O(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t} \cdot \log(\frac{1}{\epsilon}))$ (achieved using, for instance, Andrew's monotone chain convex hull algorithm).

Practical general purpose method. Due to the limited practicability of the ellipsoid algorithm, we have developed another method based on delayed column generation to solve the inner optimization problem. To simplify the presentation, we assume that $(\inf(I_r))_{r=1, \dots, R_{ij}}$ and $(\sup(I_r))_{r=1, \dots, R_{ij}}$ are all multiples of Δt . Since (12) is a linear program with a non-empty feasible set, we can equivalently compute its value by solving the dual optimization problem given by:

$$\begin{aligned} & \inf_{p_0, \dots, p_L \in \mathbb{R}} \quad \sum_{l=0, \dots, L} p_l \cdot u_j^{\Delta t}((k-l) \cdot \Delta t - \delta_{ij}^{\inf}) \\ & \text{subject to} \quad \sum_{l=0, \dots, L} p_l \cdot g_q^{ij}(l \cdot \Delta t + \delta_{ij}^{\inf}) \in [\alpha_q^{ij}, \beta_q^{ij}] \quad q = 1, \dots, Q_{ij} \\ & \quad \sum_{l=0, \dots, L} p_l = 1 \\ & \quad p_l \geq 0 \quad l = 0, \dots, L, \end{aligned} \tag{13}$$

where $L = \frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}$. Observe that the feasible set of the linear program (13) does not change across steps $k = k_i^{r,\min}, \dots, \lfloor \frac{T}{\Delta t} \rfloor$. Hence, we can warm start the primal simplex algorithm with the optimal solution found at the previous step. Furthermore, the separation oracle developed for the dual optimization problem can also be used as a subroutine for delayed column generation.

Faster procedure when the mean is the only statistics. If the ambiguity sets are only defined through a confidence interval on the mean value, i.e.:

$$\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \mathbb{E}_{X \sim p}[X] \in [\alpha^{ij}, \beta^{ij}]\},$$

then (12) can be solved to optimality in $O(\log(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}))$ computation time without resorting to the ellipsoid algorithm. First observe that (12) simplifies to:

$$\begin{aligned} & \sup_{z, y, x \in \mathbb{R}} \quad z + \alpha^{ij} \cdot x - \beta^{ij} \cdot y \\ \text{subject to} \quad & z + (x - y) \cdot l \cdot \Delta t \leq u_j^{\Delta t}((k - l) \cdot \Delta t), \quad l = \left\lceil \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rceil, \dots, \left\lfloor \frac{\delta_{ij}^{\sup}}{\Delta t} \right\rfloor \\ & z + (x - y) \cdot \delta_{ij}^{\sup} \leq u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\sup}) \\ & z + (x - y) \cdot \delta_{ij}^{\inf} \leq u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\inf}) \\ & y, x \geq 0. \end{aligned} \tag{14}$$

As it turns out, we can identify an optimal feasible basis to (14) by direct reasoning.

LEMMA 7. *An optimal solution to (14) can be found by performing three binary searches on the first coordinate of the extreme points of the upper convex hull of*

$$\{(l \cdot \Delta t, u_j^{\Delta t}((k - l) \cdot \Delta t)), l = \left\lceil \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rceil, \dots, \left\lfloor \frac{\delta_{ij}^{\sup}}{\Delta t} \right\rfloor\} \cup \{(\delta_{ij}^{\sup}, u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\sup})), (\delta_{ij}^{\inf}, u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\inf}))\}.$$

The proof is deferred to the online supplement, Section B.11.

Hence, (14) can be solved to optimality in $O(\log(\frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\Delta t}))$ running time provided that the extreme points are stored in a data structure satisfying Assumption 5.

Faster procedure when the statistics are piecewise constant. When the statistics are piecewise constant, we have:

$$a_{q,r}^{ij} = 0 \quad q = 1, \dots, Q_{ij}, r = 1, \dots, R_{ij}.$$

Hence, for any $r \in \{1, \dots, R_{ij}\}$, the set of constraints

$$z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot (a_{q,r}^{ij} \cdot l \cdot \Delta t + b_{q,r}^{ij}) \leq u_j^{\Delta t}((k - l) \cdot \Delta t) \quad l = \left\lceil \frac{\inf(I_r)}{\Delta t} \right\rceil, \dots, \left\lfloor \frac{\sup(I_r)}{\Delta t} \right\rfloor$$

is equivalent to the single constraint:

$$z + \sum_{q=1}^{Q_{ij}} (x_q - y_q) \cdot b_{q,r}^{ij} \leq \min_{l = \left\lceil \frac{\inf(I_r)}{\Delta t} \right\rceil, \dots, \left\lfloor \frac{\sup(I_r)}{\Delta t} \right\rfloor} u_j^{\Delta t}((k - l) \cdot \Delta t),$$




Figure 3 The graph of $u_j^{\Delta t}(\cdot)$ is plotted in black. The dot points represent the breakpoints of $u_j^{\Delta t}(\cdot)$.

whose right-hand side can be computed by binary search on $\mathcal{L}_{ij}^{k,r}$. As a result, the linear program (12) has $2 \cdot Q_{ij} + 1$ variables and $2 \cdot Q_{ij} + 3 \cdot R_{ij}$ constraints and can be solved to precision ϵ with an interior-point algorithm in $O(\log(\frac{1}{\epsilon}))$ computation time. Typically, piecewise constant statistics can be used to bound the probability that a given event occurs, see Section 4.3.2.

4.3.4. Dynamic convex hull algorithm Fix an arc $(i, j) \in \mathcal{A}$ and two integers $L < L'$ in $\{\left\lfloor \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rfloor, \dots, \left\lfloor \frac{\delta_{ij}^{\sup}}{\Delta t} \right\rfloor\}$. We are interested in the extreme points of the upper convex hull of $\{(l \cdot \Delta t, u_j^{\Delta t}((k-l) \cdot \Delta t)), l = L, \dots, L'\}$ for $k \in \{k_i^{r,\min}, \dots, \lfloor \frac{T}{\Delta t} \rfloor\}$. To simplify the notations, it is convenient to reverse the x-axis and shift the x-coordinate by $k \cdot \Delta t$ which leads us to equivalently look at the extreme points of the upper convex hull of:

$$\mathcal{C}_k = \{(l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t)), l = k - L', \dots, k - L\},$$

for $k \in \{k_i^{r,\min}, \dots, \lfloor \frac{T}{\Delta t} \rfloor\}$. There is a one-to-one mapping between the extreme points of these two sets which consists in applying the reverse transformation. For any k , $\hat{\mathcal{C}}_k$ denotes the upper convex hull of \mathcal{C}_k . Note that $\hat{\mathcal{C}}_k$ is a convex set and has a finitely many extreme points, all of which are in \mathcal{C}_k . Since the values $(u_j^{\Delta t}(l \cdot \Delta t))_{l=k_j^{r,\min}, \dots, \lfloor \frac{T}{\Delta t} \rfloor}$ become sequentially available in ascending order of l by chunks of size $\left\lfloor \frac{\delta_{ij}^{\inf}}{\Delta t} \right\rfloor$ as the label-setting algorithm progresses, a search for the extreme points of $\hat{\mathcal{C}}_{k+1}$ begins upon identification of the extreme points of $\hat{\mathcal{C}}_k$. Observe that $\hat{\mathcal{C}}_k$ updates to $\hat{\mathcal{C}}_{k+1}$ by removing the leftmost point $((k-L') \cdot \Delta t, u_j^{\Delta t}((k-L') \cdot \Delta t))$ and appending $((k+1-L) \cdot \Delta t, u_j^{\Delta t}((k+1-L) \cdot \Delta t))$ to the right, see Figure 3 for an illustration. In this process, deleting a point is arguably the most challenging operation because it might turn a formerly non-extreme point into one, see Figure 3b where this happens to be the case for the third leftmost point. In contrast, inserting a new point can only turn a formerly extreme point into a non-extreme one. Hence, deletions require us to do some bookkeeping other than simply keeping track of the extreme points of $\hat{\mathcal{C}}_k$ as k increases.

Maintaining the extreme points of a dynamically changing set is a well-studied class of problems




Figure 4 Local map. s and d locate the departure and arrival nodes. Three paths are highlighted. The left one (blue) is 5.3-km long and takes 9 minutes to travel. The middle one (red) is 6.4-km long and takes 8 minutes to travel. The rightmost one (green) is 6.1-km long and takes 10 minutes to travel.

in computational geometry known as *Dynamic Convex Hull* problems. Specific instances from this class differ along the operations to be performed on the set (e.g. insertions, deletions), the queries to be answered on the extreme points, and the dimensionality of the input data. Brodal and Jacob (2002) design a data structure maintaining a description of the upper convex hull of a finite set of N points in \mathbb{R}^2 . This data structure satisfies Assumption 5 as it allows to insert points, to delete points, and to perform binary search on the first coordinate of the extreme points, all in amortized time $O(\log(N))$ and with $O(N)$ space usage. For the purpose of being self-contained, we design our own data structure in the online supplement Section A to tackle the particular dynamic convex hull problem at hand. Our approach is based on Andrew’s monotone chain convex hull algorithm, see Andrew (1979), and only uses two arrays and a stack. The data structure developed in Brodal and Jacob (2002) is more complex than ours but can handle arbitrary dynamic convex hull problems.

5. Numerical experiments

In this section, we compare, using a real-world application with field data from the Singapore road network, the performance of the nominal and robust approaches to vehicle routing when traffic measurements are scarce and uncertain. To benchmark the performance of the robust approach, we propose a realistic framework where both the nominal and robust approaches can be efficiently computed and for which it is up to the user to pick one.

5.1. Framework

We work on a network composed of the main roads of Singapore with 20,221 arcs and 11,018 nodes for a total length of 1131 kilometers of roads. The data consists of a 15-day recording of GPS probe vehicle speed samples coming from a combined fleet of over 15,000 taxis. Features of each recording include current location, speed and status (free, waiting for a customer, occupied). We denote by s and d the departure and arrival nodes. Because there is usually only one reasonable route to get from s to d for most pairs (s, d) in our network, the benefits of using one vehicle routing approach over

another would not be apparent if we were to pick (s, d) uniformly at random over \mathcal{V}^2 . Instead, we choose to hand-pick a pair (s, d) with at least two reasonable routes to get from s to d with similar travel times so that the best driving itinerary depends on the actual traffic conditions. We choose $s = \text{“Woodlands avenue 2”}$ and $d = \text{“Mandai link”}$, see Figure 4, but the results would be similar for other pairs satisfying this property.

Method of performance evaluation. Consider the following real-world situation. A user has to find an itinerary to get from s to d within a given budget T (the deadline) and with an objective to maximize the probability of on-time arrival, but when only a few vehicle speed samples are available in order to assess arc travel time uncertainty.

To model this real-world situation, we assume that the full set of samples of vehicle speed measurement available in our dataset in fact represents the real traffic conditions, characterized by the corresponding travel-time distributions p_{ij}^{real} ’s, which are obtained from the full set of samples. Mimicking the fact that the p_{ij}^{real} ’s are actually not fully available, we then consider the case where only a fraction of the full set of samples, say $\lambda \in [0, 1]$, is available. Based on this limited data, the challenge is to select an itinerary with a probability of on-time arrival with respect to the real traffic conditions p_{ij}^{real} ’s as high as possible. We propose to use the methods listed in Table 2 to choose such an itinerary. For each of these methods, the process goes as follows:

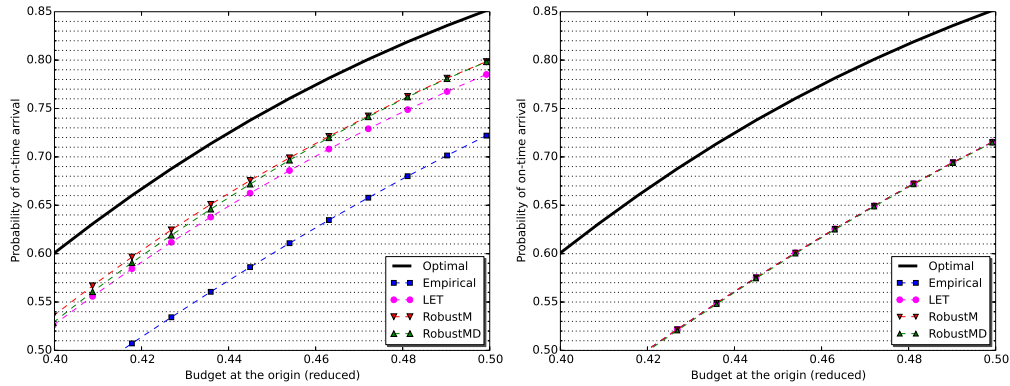
1. Estimate the arc-based travel-time parameters required to run the method using the fraction of data available.
2. Run the corresponding algorithm to find an itinerary, depending on the chosen method.
3. Compute the probability of on-time arrival of this itinerary for the **real** traffic conditions ($\lambda = 1$).

The result obtained depends on both λ and the available samples as there are many ways to pick a fraction λ out of the entire dataset. Hence, for each λ in a set Λ , we randomly pick $\lambda \cdot N_{ij}$ samples for each arc (i, j) , where N_{ij} is the number of samples collected in the entire dataset for that particular arc. For each $\lambda \in \Lambda$, and for each method, we store the calculated probability of on-time arrival. We repeat this procedure 100 times.

A few remarks are in order. We choose $\Lambda = \{0.001, 0.002, 0.005\}$, this corresponds to an average number of samples per arc of $[5.5, 9.4, 25.1]$ respectively (we take at least one sample per arc). The average arc length is 163 meters, hence we set $\Delta t = 0.02$ second to get a good accuracy. This parameter has a significant impact on the running time and it could also be optimized. We include the LET method as it is a reasonably robust approach, although not tailored to the risk function considered, and because it is very fast to solve. The confidence intervals used by the robust approaches are percentile bootstrap 95 % confidence intervals derived from resampling the available data with replacement. When solving the discretization schemes (6) and (9), ties in the argument of the maximum are broken in favor of the (estimated) least expected travel time to the destination. To solve the

Table 2 Methods considered. I_{ij}^m and I_{ij}^{md} are confidence intervals.

Method	Travel-time parameters to estimate from samples	Approach
RobustM	$\delta_{ij}^{\inf}, \delta_{ij}^{\sup}, I_{ij}^m$	(3) with $\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \mathbb{E}_{X \sim p}[X] \in I_{ij}^m\}$
RobustMD	$\delta_{ij}^{\inf}, \delta_{ij}^{\sup}, I_{ij}^m, I_{ij}^{md},$ $m_{ij} = \frac{\max(I_{ij}^m) + \min(I_{ij}^m)}{2}$	(3) with $\mathcal{P}_{ij} = \{p \in \mathcal{P}([\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]) : \begin{matrix} \mathbb{E}_{X \sim p}[X] \in I_{ij}^m \\ \mathbb{E}_{X \sim p}[X - m_{ij}] \in I_{ij}^{md} \end{matrix}\}$
Empirical	empirical distributions p_{ij}	(1) with p_{ij}
LET	empirical mean m_{ij}	standard shortest path



(a) Average probability of on-time arrival. (b) 5% worst-case probability of on-time arrival.

Figure 5 $\lambda = 0.001$, average number of samples per link: ~ 5.5 .

robust problems, we use the column generation scheme and the special-purpose procedure described in Section 4.3.3 while we use the scheme based on fast Fourier transforms described in Section 3.2.2 for the nominal approach.

5.2. Results

The results are plotted in Figure 5, 6, and 7. Each of these figures corresponds to one of the fraction $\lambda \in \Lambda$ so as to see the impact of an increasing knowledge. The time budget is “normalized”: 0 (resp. 1) corresponds to the minimum (resp. maximum) amount of time it takes to reach d from s . For each λ , for each method in Table 2, for each time budget T , and for each of the 100 simulations, we compute the actual probability of on-time arrival of the corresponding strategy. The average (resp. 5 % worst-case) probability of on-time arrival over the simulations is plotted on the figures labeled “a” (resp. “b”). The 5 % worst-case measure, which corresponds to the average over the 5 simulations out 100 that yield the lowest probability of arriving on-time, is particularly relevant as commuters opting for this risk function would expect the approach to have good results even under bad scenarios. We also plot the average runtime for each of the method as a function of the time budget in Figure 8.

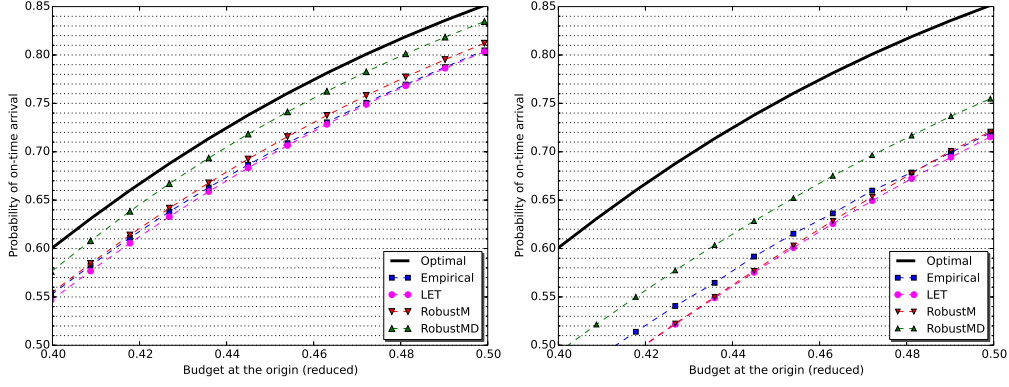


Figure 6 (a) Average probability of on-time arrival. (b) 5% worst-case probability of on-time arrival. $\lambda = 0.002$, average number of samples per link: ~ 9.4 .

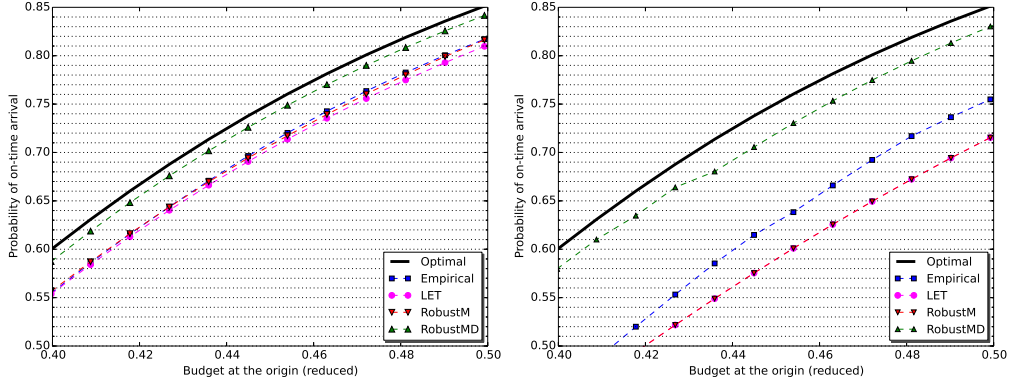


Figure 7 (a) Average probability of on-time arrival. (b) 5% worst-case probability of on-time arrival. $\lambda = 0.005$, average number of samples per link: ~ 25.1 .

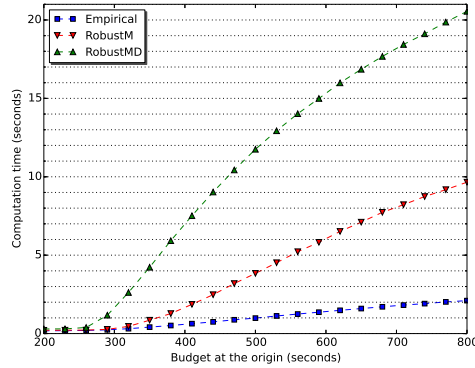


Figure 8 Average computation time as a function of the time budget for $\lambda = 0.001$.

Conclusions. As can be observed on the figures, Empirical is not competitive when only a few samples are available. To be specific, RobustM slightly outperforms the other methods when there are very few measurements, see Figure 5, while RobustMD is a clear winner when more samples are

available, in terms of both average and worst-case performances, see Figures 6 and 7. Observe that, as expected, the performance of Empirical improves as more samples get available and Empirical eventually outperforms RobustM, see Figure 7. Our interpretation of these results is that relying on quantities, either moments or distributions, that cannot be accurately estimated may be misleading even for robust strategies. On the other hand, failure to capture the increasing knowledge on the actual travel-time probability distributions (e.g. by estimating more moments) as the amount of available data increases may lead to poor performances.

6. Extensions

In this section, we sketch how to extend the results derived in Sections 3 and 4 when either Assumption 1 or Assumption 2 is relaxed. Most of the results also extend when both assumptions are relaxed at the same time but we choose to discuss one assumption at a time to highlight their respective implications.

6.1. Relaxing the independence Assumption 1: Markovian costs

We consider here the case where the experienced costs of crossing arcs define a *Markov* chain of finite order m . To simplify the presentation, we provide in details the extensions of our previous results to the case $m = 1$. Adapting these extensions to a general m amounts to augmenting the state space of the underlying MDP by the costs of the last m visited arcs. We emphasize that while *Markov* chains can model the reality of the decision making process more accurately, this comes at a price: this requires an estimation of m -dimensional probability distributions, and the computational time needed to find an optimal strategy grows exponentially with m .

Extension for the nominal problem. A variant of Theorem 1 can be shown to hold if the arc cost distributions are discrete. Under this assumption and as soon as the total cost spent so far is larger than $T - T_f$, the optimal strategy coincides with the strategy of minimizing the expected costs, which may no longer be a shortest path but can still be shown to be a solution without cycles. Under the same assumption, Proposition 1 remains valid under the following higher-dimensional dynamic program:

$$\begin{aligned}
 u_d(t, z, \theta) &= f(t) & t \leq T, z \in \mathcal{A}(d), \theta \in \Theta_{zd} \\
 u_i(t, z, \theta) &= \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega | z, \theta) \cdot u_j(t - \omega, i, \omega) d\omega & i \neq d, t \leq T, z \in \mathcal{A}(i), \theta \in \Theta_{zi} \\
 \pi_f^*(i, t, z, \theta) &\in \arg \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}(\omega | z, \theta) \cdot u_j(t - \omega, i, \omega) d\omega & i \neq d, t \leq T, z \in \mathcal{A}(i), \theta \in \Theta_{zi},
 \end{aligned} \tag{15}$$

where $\mathcal{A}(i)$ denotes the set of immediate antecedents of i in \mathcal{G} , Θ_{zi} is the finite set of possible values taken by c_{zi} for $z \in \mathcal{A}(i)$, and $p_{ij}(\cdot | z, \theta)$ is the conditional distribution of c_{ij} given that the last visited node is z and that $c_{zi} = \theta$. The discretization scheme of Section 3.2.1 can be adapted for this

new dynamic equation and the approximation guarantees carry over. To solve this new discretization scheme, the label-setting approach from Section 3.2.2 can be adapted by observing that the functions $(u_i(\cdot, z, \theta))_{i \in \mathcal{V}, z \in \mathcal{A}(i), \theta \in \Theta_{zi}}$ can be computed block by block by interval increments of size δ^{inf} . However, the schemes based on fast Fourier transforms and the idea of zero-delay convolution do not apply anymore, and we need to use the pointwise definition of convolution products with computational complexity:

$$O\left(\max_{(i,j) \in \mathcal{A}} |\Theta_{ij}| \cdot \frac{|\mathcal{A}| \cdot (T - T_f) + |\mathcal{V}|^2 \cdot \delta^{\text{sup}}}{\Delta t}\right).$$

Extension for the robust problem. For any $(i, j) \in \mathcal{A}$, $z \in \mathcal{A}(i)$, and $\theta \in \Theta_{zi}$, $p_{ij}(\cdot | z, \theta)$ is only known to lie in the ambiguity set $\mathcal{P}_{ij,z,\theta}$. If $\mathcal{P}_{ij,z,\theta}$ is only comprised of discrete distributions with finite support Θ_{ij} , a variant of Theorem 2 can be shown to hold. Specifically, as soon as the total cost spent so far is larger than $T - T_f^r$, the optimal strategy coincides with the strategy of minimizing the worst-case expected costs, which can also be shown not to cycle. Under this assumption, Proposition 3 remains valid under the following higher-dimensional dynamic program:

$$\begin{aligned} u_d(t, z, \theta) &= f(t) & t \leq T, z \in \mathcal{A}(d), \theta \in \Theta_{zd} \\ u_i(t, z, \theta) &= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij,z,\theta}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega, i, \omega) d\omega & i \neq d, t \leq T, z \in \mathcal{A}(i), \theta \in \Theta_{zi} \\ \pi_{f,\mathcal{P}}^*(i, t, z, \theta) &\in \arg \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij,z,\theta}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t - \omega, i, \omega) d\omega & i \neq d, t \leq T, z \in \mathcal{A}(i), \theta \in \Theta_{zi}. \end{aligned} \quad (16)$$

The discretization scheme of Section 4.3.1 can be adapted for this new set of equations and the approximation guarantees of Proposition 4 carry over. Moreover, the label-setting approach can also be adapted along the same lines as for the nominal problem. The ideas underlying the algorithmic developments of Section 4.3.3 remain valid but we now have to recompute the convex hulls from scratch at each time step using Andrew's monotone chain convex hull algorithm, as opposed to using a dynamic convex hull algorithm, which leads to the computational complexity:

$$O\left(\max_{(i,j) \in \mathcal{A}} |\Theta_{ij}| \cdot \log\left(\max_{(i,j) \in \mathcal{A}} |\Theta_{ij}|\right) \cdot \frac{|\mathcal{A}| \cdot (T - T_f^r) + |\mathcal{V}|^2 \cdot \delta^{\text{sup}}}{\Delta t} \cdot \log\left(\frac{|\mathcal{V}| + \frac{T - T_f^r}{\delta^{\text{inf}}}}{\epsilon}\right)\right),$$

when we want to compute an ϵ -approximate strategy solution to the discretization scheme (9).

6.2. Relaxing Assumption 2: τ -dependent arc cost probability distributions

Extension for the nominal problem. For any $\tau \geq 0$ and $(i, j) \in \mathcal{A}$, we denote by p_{ij}^τ the distribution of c_{ij}^τ and by m_{ij}^τ the mean of p_{ij}^τ . Theorem 1 remains valid if, for any $(i, j) \in \mathcal{A}$, m_{ij}^τ converges as $\tau \rightarrow \infty$, in which case the shortest-path tree mentioned in the statement is defined with respect to the limits of the mean arc costs. For instance, this assumption is satisfied when the distributions are time-varying during a peak period and stationary anytime thereafter, see Miller-Hooks and Mahmassani

(2000). Under this assumption, Proposition 1 also remains valid but for the slightly modified dynamic program:

$$\begin{aligned} u_d(t) &= f(t) & t \leq T \\ u_i(t) &= \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}^{T-t}(\omega) \cdot u_j(t-\omega) d\omega & i \neq d, t \leq T \\ \pi_f^*(i, t) &\in \arg \max_{j \in \mathcal{V}(i)} \int_0^\infty p_{ij}^{T-t}(\omega) \cdot u_j(t-\omega) d\omega & i \neq d, t \leq T. \end{aligned} \quad (17)$$

The discretization scheme of Section 3.2.1 can be trivially adapted for this new dynamic equation, although we may loose the approximation guarantees provided by Proposition 2. For them to carry over, we need additional assumptions. To be specific, one of the following properties must be satisfied:

- the arc cost distributions vary smoothly, in the sense that, for any arc $(i, j) \in \mathcal{A}$, there exists K such that the Kolmogorov distance between $p_{ij}^{\tau_1}$ and $p_{ij}^{\tau_2}$ is smaller than $K \cdot |\tau_1 - \tau_2|$ for any $\tau_1, \tau_2 \geq 0$,
- the arc cost distributions are discrete and the discretization length Δt is chosen appropriately,
- the arc cost distributions change finitely many times and the discretization length Δt is chosen appropriately.

To solve the discretization scheme, the label-setting approach described in Section 3.2.2 remains relevant but we now have to apply the pointwise definition of convolution products, as opposed to using fast Fourier transforms and zero-delay convolutions, with computational complexity quadratic in $\frac{1}{\Delta t}$:

$$O\left(\frac{|\mathcal{A}| \cdot (T - T_f) + |\mathcal{V}|^2 \cdot \delta^{\sup}}{\Delta t} \cdot \frac{\delta^{\sup} - \delta^{\inf}}{\Delta t}\right).$$

Extension for the robust problem. For any $\tau \geq 0$ and $(i, j) \in \mathcal{A}$, p_{ij}^τ is only known to lie in the ambiguity set \mathcal{P}_{ij}^τ . First observe that (3) turns into:

$$\sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i, j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}^\tau} \mathbb{E}_{\mathbf{P}^\tau}[f(T - X_\pi)],$$

which is exactly the robust counterpart of (1), as opposed to a robust relaxation when the arc cost distributions are stationary. Theorem 2 remains valid if, for any $(i, j) \in \mathcal{A}$, $\max_{p_{ij} \in \mathcal{P}_{ij}^\tau} \mathbb{E}_{X \sim p_{ij}}[X]$ converges as $\tau \rightarrow \infty$, in which case the shortest-path tree mentioned in the statement is defined with respect to the limits. Again, this assumption is, for instance, satisfied when the ambiguity sets are time-varying during a peak period and stationary anytime thereafter. Under this assumption, Proposition 3 also remains valid but for the slightly modified dynamic program:

$$\begin{aligned} u_d(t) &= f(t) & t \leq T \\ u_i(t) &= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}^{T-t}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t-\omega) d\omega & i \neq d, t \leq T \\ \pi_{f, \mathcal{P}}^*(i, t) &\in \arg \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}^{T-t}} \int_0^\infty p_{ij}(\omega) \cdot u_j(t-\omega) d\omega & i \neq d, t \leq T. \end{aligned} \quad (18)$$

Similarly as for the nominal problem, the discretization scheme can be trivially adapted but we may loose the approximation guarantees provided by Proposition 4. For them to carry over, one of the following properties has to be satisfied:

- the ambiguity sets vary smoothly, in the sense that, for any arc $(i, j) \in \mathcal{A}$, there exists K such that the Kolmogorov distance between $\mathcal{P}_{ij}^{\tau_1}$ and $\mathcal{P}_{ij}^{\tau_2}$ is smaller than $K \cdot |\tau_1 - \tau_2|$ for any $\tau_1, \tau_2 \geq 0$,
- the ambiguity sets are only comprised of discrete distributions and the discretization length Δt is chosen appropriately,
- the ambiguity sets change finitely many times and the discretization length Δt is chosen appropriately.

In contrast to the nominal problem, all the algorithms developed in Section 4.3.3 can still be used to solve the discretization scheme with the same computational complexity as long as the ambiguity sets are defined by confidence intervals on piecewise affine statistics, as precisely defined in Section 4.3.2.

Acknowledgments

This research is supported in part by the National Research Foundation (NRF) Singapore through the Singapore MIT Alliance for Research and Technology (SMART) and its Future Urban Mobility (FM) Interdisciplinary Research Group. The authors would like to thank Chong Yang Goh from the Massachusetts Institute of Technology for his help in preprocessing the data.

References

- Adulyasak, Y., P. Jaillet. 2014. Models and algorithms for stochastic and robust vehicle routing with deadlines. *Transportation Sci.* (Articles in Advance).
- Andrew, A. 1979. Another efficient algorithm for convex hulls in two dimensions. *Inform. Processing Lett.* **9**(5) 216–219.
- Bertsekas, D. P., J. Tsitsiklis. 1991. An analysis of stochastic shortest path problems. *Math. Oper. Res.* **16**(3) 580–595.
- Bertsimas, D., I. Popescu. 2005. Optimal inequalities in probability theory: A convex optimization approach. *SIAM J. Optim.* **15**(3) 780–804.
- Brodal, G. S., R. Jacob. 2002. Dynamic planar convex hull. *Proc. 43rd IEEE Annual Symp. Foundations Comput. Sci.*. IEEE, 617–626.
- Calafiore, C., L. El Ghaoui. 2006. On distributionally robust chance-constrained linear programs. *J. Optim. Theory and Applications* **130**(1) 1–22.
- Dean, B. 2010. Speeding up stochastic dynamic programming with zero-delay convolution. *Algorithmic Oper. Res.* **5**(2) 96–104.

- Delage, E., Y. Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* **58**(3) 595–612.
- Fan, Y., R. Kalaba, I. Moore. 2005. Arriving on time. *J. Optim. Theory and Applications* **127**(3) 497–513.
- Frank, H. 1969. Shortest paths in probabilistic graphs. *Oper. Res.* **17**(4) 583–599.
- Gabrel, V., C. Murat, L. Wu. 2013. New models for the robust shortest path problem: complexity, resolution and generalization. *Annals Oper. Res.* **207** 97–120.
- Hoy, D., E. Nikolova. 2015. Approximately optimal risk-averse routing policies via adaptive discretization. *Proc. 29th Internat. Conf. Artificial Intelligence (AAAI)*.
- Iyengar, G. 2005. Robust dynamic programming. *Math. of Oper. Res.* **30**(2) 257–280.
- Jaillet, P., J. Qi, M. Sim. 2015. Routing optimization with deadlines under uncertainty. *Oper. Res.* Forthcoming.
- Loui, R. P. 1983. Optimal paths in graphs with stochastic or multidimensional weights. *Comm. ACM* **26**(9) 670–676.
- Miller-Hooks, E., H. Mahmassani. 2000. Least expected time paths in stochastic, time-varying transportation networks. *Transportation Sci.* **34**(2) 198–215.
- Nie, Y., Y. Fan. 2006. Arriving-on-time problem: discrete algorithm that ensures convergence. *Transportation Res. Record* **1964** 193–200.
- Nie, Y., X. Wu. 2009. Shortest path problem considering on-time arrival probability. *Transportation Res. B* **43**(6) 597–613.
- Nikolova, E., M. Brand, D. R. Karger. 2006a. Optimal route planning under uncertainty. *Proc. Internat. Conf. Automated Planning Scheduling*. 131–140.
- Nikolova, E., J. A. Kelner, M. Brand, M. Mitzenmacher. 2006b. Stochastic shortest paths via quasi-convex maximization. *Proc. 14th Annual Eur. Sympos. Algorithms (ESA '06)*, vol. 14. Springer Berlin Heidelberg, 552–563.
- Nilim, A., L. El Ghaoui. 2005. Robust control of markov decision processes with uncertain transition matrices. *Oper. Res.* **53**(5) 780–798.
- Overmars, M. H., J. Van Leeuwen. 1981. Maintenance of configurations in the plane. *J. Comput. and System Sci.* **23**(2) 166–204.
- Parmentier, A., F. Meunier. 2014. Stochastic shortest paths and risk measures. *arXiv preprint arXiv:1408.0272*.
- Prékopa, A. 1990. The discrete moment problem and linear programming. *Discrete Applied Math.* **27**(3) 235–254.
- Puterman, M. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

- Samaranayake, S., S. Blandin, A. Bayen. 2012a. Speedup techniques for the stochastic on-time arrival problem. *12th Workshop Algorithmic Approaches Transportation Model. Optim. Systems (ATMOS 2012)*, vol. 25. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 83–96.
- Samaranayake, S., S. Blandin, A. Bayen. 2012b. A tractable class of algorithms for reliable routing in stochastic networks. *Transportation Res. C* **20**(1) 199–217.
- Shapiro, A. 2001. On duality theory of conic linear problems. *Semi-Infinite Programming*, vol. 57. Kluwer Academic Publishers, 135–165.
- Vandenbergh, L., S. Boyd, K. Comanor. 2007. Generalized Chebyshev bounds via semidefinite programming. *SIAM Rev.* **49**(1) 52–64.
- White III, C. C., H. K. Eldeib. 1994. Markov decision processes with imprecise transition probabilities. *Oper. Res.* **42**(4) 739–749.
- Wiesemann, W., D. Kuhn, B. Rustem. 2013. Robust markov decision processes. *Math. of Oper. Res.* **38**(1) 153–183.
- Wiesemann, W., D. Kuhn, M. Sim. 2014. Distributionally robust convex optimization. *Oper. Res.* **62**(6) 1358–1376.
- Xu, H., C. Caramanis, S. Mannor. 2012. Optimization under probabilistic envelope constraints. *Oper. Res.* **60**(3) 682–699.
- Xu, H., S. Mannor. 2010. Distributionally robust markov decision processes. *Adv. Neural Inform. Processing Systems*. 2505–2513.
- Xu, H., S. Mannor. 2011. Probabilistic goal Markov decision processes. *Proc. 22th Internat. Joint Conf. Artificial Intelligence (AAAI)*. AAAI Press, 2046–2052.

Online Supplement

Appendix A: Tailored dynamic convex hull algorithm

The fact that deletions and insertions always occur on the same side of the set allows us to deal with deletions in an indirect way, by building and merging upper convex hulls of partial input data. The only downside is that this requires an efficient merging procedure. In this respect, we state without proof a result derived from Overmars and Leeuwen (1981).

LEMMA EC.1. *Consider a set S of N points in \mathbb{R}^2 partitioned into two sets of points S_1 and S_2 such that, for any two points $(x_1, y_1) \in S_1$ and $(x_2, y_2) \in S_2$ we have $x_1 < x_2$. Suppose that the extreme points of \hat{S}_1 (resp. \hat{S}_2) are stored in an array A_1 (resp. A_2) of size N in ascending order of their first coordinates. We can find two indices l_1 and l_2 in $O(\log(N))$ time such that the set comprised of the points contained in A_1 with index smaller than l_1 and the points contained in A_2 with index larger than l_2 is precisely the set of extreme points of \hat{S} .*

Algorithm. We use two arrays A_{left} and A_{right} along with a stack \mathcal{S} . The arrays A_{left} and A_{right} are of size $L' - L + 1$, indexed from 0 to $L' - L$, and store points in \mathbb{R}^2 in ascending order of their first coordinates. The stack \mathcal{S} stores stacks of points in \mathbb{R}^2 . We keep track of two indices l_{left} and l_{right} such that, at any step $k = k_i^{r, \min} + p \cdot (L' - L + 2) + r$ for some $p \in \mathbb{N}$ and $0 \leq r \leq L' - L + 1$, the following invariant holds:

- $\{A_{\text{left}}[l], l = l_{\text{left}} + 1, \dots, L' - L\}$ is the set of extreme points of the upper convex hull of $\{(l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t)), l = k - L', \dots, k - L - r\}$,
- $\{A_{\text{right}}[l], l = 0, \dots, l_{\text{right}} - 1\}$ is the set of extreme points of the upper convex hull of $\{(l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t)), l = k - L - r + 1, \dots, k - L\}$.

Using the procedure of Lemma EC.1 and this invariant, we can find a pair of indices (l_1, l_2) in $O(\log(L' - L))$ time such that $\{A_{\text{left}}[l], l = l_{\text{left}} + 1, \dots, l_1\} \cup \{A_{\text{right}}[l], l = l_2, \dots, l_{\text{right}} - 1\}$ is the set of extreme points of $\hat{\mathcal{C}}_k$. Hence, all we have left to do is to provide a procedure to maintain A_{left} , A_{right} , l_{left} and l_{right} , which we do next.

A_{left} , A_{right} , and \mathcal{S} are initially empty. The algorithm proceeds in two phases and loops back to the first one every $L' - L + 2$ steps. For convenience, we define cross as the function taking as an input three points a, b, c in \mathbb{R}^2 and returning the cross product of the vector \vec{ab} and \vec{ac} .

Phase 1: Suppose that the current step is k . Hence, the values $u_j^{\Delta t}(k - L'), \dots, u_j^{\Delta t}(k - L)$ are available. This phase is based on Andrew's monotone chain convex hull algorithm to find the extreme points of $\hat{\mathcal{C}}_k$ with the difference that we store the points removed along the process in stacks for future use. Specifically, set $l_{\text{left}} = L' - L$ and $l_{\text{right}} = 0$ and for l decreasing from $k - L$ to $k - L'$, do the following:

- (a) Initialize a new stack \mathcal{S}' ,
- (b) While $l_{\text{left}} \leq L' - L - 2$ and $\text{cross}(A_{\text{left}}[l_{\text{left}} + 2], A_{\text{left}}[l_{\text{left}} + 1], (l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t))) \geq 0$:
 - Push $A_{\text{left}}[l_{\text{left}} + 1]$ to \mathcal{S}' ,
 - Increment l_{left} ,
- (c) Push \mathcal{S}' to \mathcal{S} ,
- (d) Set $A_{\text{left}}[l_{\text{left}}] = (l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t))$ and decrement l_{left} . At this point, $\{A_{\text{left}}[l], l = l_{\text{left}} + 1, \dots, L' - L\}$ is the set of extreme points of the upper convex hull of $\{(m \cdot \Delta t, u_j^{\Delta t}(m \cdot \Delta t)), m = l, \dots, k - L\}$.

Phase 2: At step $k + l$, for l increasing from 1 to $L' - L$, observe that the value $u_j^{\Delta t}(k + l - L)$ becomes available. To maintain A_{left} and l_{left} , we remove the leftmost point (x, y) and reinsert the points, stored in the topmost stack of \mathcal{S} , that were previously removed from A_{left} when appending (x, y) to A_{left} in the course of running Andrew's monotone chain convex hull algorithm. Specifically:

- (a) Increment l_{left} ,
- (b) Pop the topmost stack \mathcal{S}' out of \mathcal{S} ,
- (c) While \mathcal{S}' is not empty:
 - Pop the topmost point (x, y) of \mathcal{S}' ,
 - Set $A_{\text{left}}[l_{\text{left}}] = (x, y)$,
 - Decrement l_{left} .

To maintain A_{right} and l_{right} , we run an iteration of Andrew's monotone chain convex hull algorithm. Specifically:

- (a) While $l_{\text{right}} \geq 2$ and $\text{cross}(A_{\text{right}}[l_{\text{right}} - 2], A_{\text{right}}[l_{\text{right}} - 1], ((k + l - L) \cdot \Delta t, u_j^{\Delta t}((k + l - L) \cdot \Delta t))) \leq 0$:
 - Decrement l_{right} ,
- (b) Set $A_{\text{right}}[l_{\text{right}}] = ((k + l - L) \cdot \Delta t, u_j^{\Delta t}((k + l - L) \cdot \Delta t))$ and increment l_{right} .

Complexity Analysis. Observe that any point added to A_{left} can only be removed once, and the same holds for A_{right} . This means that Phase 1 and Phase 2 take $O(L' - L)$ computation time. These two phases are repeated $\left\lceil \frac{\lfloor \frac{T}{\Delta t} \rfloor - k_i^{r, \min}}{L' - L + 2} \right\rceil$ times leading to an overall complexity of $O(\lfloor \frac{T}{\Delta t} \rfloor - k_i^{r, \min})$. Since the merging procedure outlined in Lemma EC.1 takes $O(\log(L' - L))$ computation time at each step, the global complexity is $O((\lfloor \frac{T}{\Delta t} \rfloor - k_i^{r, \min}) \cdot \log(L' - L))$.

Appendix B: Omitted Proofs

B.1. Proof of Theorem 1

Proof of Theorem 1. Let us start with the last part of the theorem. If the support of $f(\cdot)$ is included in $[T_f, \infty)$, any strategy is optimal when having already spent a budget of $T - T_f$ with an optimal objective function of 0.

Let us now focus on the first part of the theorem. Consider an optimal strategy π_f^* solution to

(1). For a given history $h \in \mathcal{H}$, we define t_h as the remaining budget, i.e. T minus the total cost spent so far, and i_h as the current location. The policy π_f^* maps $h \in \mathcal{H}$ to a probability distribution over $\mathcal{V}(i_h)$. Observe that randomizing does not help because the costs are independent across time and arcs so that, without loss of generality, we can assume that π_f^* actually maps h to the node in $\mathcal{V}(i_h)$ minimizing the objective function given h . For $h \in \mathcal{H}$, we denote by $X_{\pi_f^*}^h$ the random cost-to-go incurred by following strategy π_f^* , i.e. not including the total cost spent up to this point of the history $T - t_h$. We define $(m_{ij})_{(i,j) \in \mathcal{A}}$ as the expected arc costs and by M_i as the minimum expected cost to go from i to d for any $i \in \mathcal{V}$. We also define π_s as a policy associated with an arbitrary shortest path from i to d with respect to the expected costs. Specifically, π_s maps the current location i_h to a node in $\mathcal{T}(i_h)$, irrespective of the history of the process. Similarly as for π_f^* , we denote by $X_{\pi_s}^h$ the random cost-to-go incurred by following strategy π_s for $h \in \mathcal{H}$. We first show that there exists T_f such that, for both cases (a) and (b):

$$\mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} < \min_{i \neq d} \min_{j \in \mathcal{V}(i), j \notin \mathcal{T}(i)} \{m_{ij} + M_j - M_i\} \quad \forall h \in \mathcal{H} \text{ such that } t_h \leq T_f, \quad (\text{EC.1})$$

with the convention that the minimum of an empty set is equal to infinity. Note that the right-hand side is always positive. Let $\alpha = |\mathcal{V}| \cdot \delta^{\sup}$.

(a) For $h \in \mathcal{H}$ such that $t_h < T_1$, we have, using a Taylor's series expansion:

$$f(t_h - X_{\pi_f^*}^h) = f(t_h - \alpha) + f'(t_h - \alpha) \cdot (\alpha - X_{\pi_f^*}^h) + \frac{1}{2} \cdot f''(\xi_h) \cdot (\alpha - X_{\pi_f^*}^h)^2,$$

where $\xi_h \in [\min(t_h - \alpha, t_h - X_{\pi_f^*}^h), \max(t_h - \alpha, t_h - X_{\pi_f^*}^h)]$, and:

$$f(t_h - X_{\pi_s}^h) = f(t_h - \alpha) + f'(t_h - \alpha) \cdot (\alpha - X_{\pi_s}^h) + \frac{1}{2} \cdot f''(\zeta_h) \cdot (\alpha - X_{\pi_s}^h)^2,$$

where $\zeta_h \in [\min(t_h - \alpha, t_h - X_{\pi_s}^h), \max(t_h - \alpha, t_h - X_{\pi_s}^h)]$. Using *Bellman's Principle of Optimality* for π_f^* , we have:

$$\mathbb{E}[f(t_h - X_{\pi_f^*}^h)] = \mathbb{E}[f(T - ((T - t_h) + X_{\pi_f^*}^h))] \geq \mathbb{E}[f(T - ((T - t_h) + X_{\pi_s}^h))] \geq \mathbb{E}[f(t_h - X_{\pi_s}^h)].$$

Expanding and rearranging yields:

$$-f'(t_h - \alpha) \cdot (\mathbb{E}[X_{\pi_f^*}^h] - \mathbb{E}[X_{\pi_s}^h]) \geq \frac{1}{2} \cdot (\mathbb{E}[-f''(\xi_h) \cdot (\alpha - X_{\pi_f^*}^h)^2] + \mathbb{E}[f''(\zeta_h) \cdot (\alpha - X_{\pi_s}^h)^2]).$$

Since the costs are independent across time and arcs:

$$\mathbb{E}[X_{\pi_s}^h] = M_{i_h}.$$

Concavity of $f(\cdot)$ implies that $f''(\xi_h) \cdot (\alpha - X_{\pi_f^*}^h)^2 \leq 0$ almost surely. Since $f(\cdot)_{(-\infty, T_1)}$ is increasing, we obtain $\mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} \leq \frac{\mathbb{E}[-f''(\zeta_h) \cdot (\alpha - X_{\pi_s}^h)^2]}{2 \cdot f'(t_h - \alpha)}$. As $X_{\pi_s}^h$ is the cost of a path, Assumption 3 implies $0 \leq X_{\pi_s}^h \leq \alpha$. We get that $\zeta_h \in [t_h - \alpha, t_h]$ and:

$$\mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} \leq -\alpha^2 \cdot \frac{\inf_{[t_h - \alpha, t_h]} f''}{2 \cdot f'(t_h - \alpha)}.$$

As $f''(\cdot)$ is continuous, there exists $\alpha_{t_h} \in [0, \alpha]$ such that $\inf_{[t_h - \alpha, t_h]} f'' = f''(t_h - \alpha_{t_h})$. Since $f'(\cdot)$ is non-increasing on $(-\infty, T_1)$, we derive:

$$\mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} \leq -\alpha^2 \cdot \frac{f''(t_h - \alpha_{t_h})}{2 \cdot f'(t_h - \alpha_{t_h})}.$$

By assumption $\frac{f''}{f'}(\cdot)$ vanishes at $-\infty$ therefore we can pick T_f small enough to get the desired inequality.

(b) As $f' \rightarrow_{-\infty} a > 0$, we can find $T_f < T_1$ small enough such that:

$$|f'(t) - a| < \epsilon \quad \forall t \leq T_f,$$

with $\epsilon = a \cdot \frac{\beta}{2\alpha + \beta}$ and where β is the right-hand side of the desired inequality. Consider $h \in \mathcal{H}$ such that $t_h \leq T_f$. Using *Bellman's Principle of Optimality* for π_f^* , we have:

$$\mathbb{E}[f(t_h - X_{\pi_f^*}^h)] = \mathbb{E}[f(T - ((T - t_h) + X_{\pi_f^*}^h))] \geq \mathbb{E}[f(T - ((T - t_h) + X_{\pi_s}^h))] \geq \mathbb{E}[f(t_h - X_{\pi_s}^h)].$$

Since f is C^1 on $(-\infty, T_f)$, this yields:

$$\begin{aligned} 0 &\leq \mathbb{E}[f(t_h - X_{\pi_f^*}^h) - f(t_h - X_{\pi_s}^h)] \\ &\leq \mathbb{E}[f(t_h - X_{\pi_f^*}^h) - f(t_h) + f(t_h) - f(t_h - X_{\pi_s}^h)] \\ &\leq \mathbb{E}\left[-\int_{t_h - X_{\pi_f^*}^h}^{t_h} f' + \int_{t_h - X_{\pi_s}^h}^{t_h} f'\right] \\ &\leq \mathbb{E}[-(a - \epsilon) \cdot X_{\pi_f^*}^h + (a + \epsilon) \cdot X_{\pi_s}^h]. \end{aligned}$$

Since the costs are independent across time and arcs:

$$\mathbb{E}[X_{\pi_s}^h] = M_{i_h}.$$

Rearranging the last inequality, we derive:

$$\begin{aligned} \mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} &\leq \frac{2\epsilon}{a - \epsilon} \cdot M_{i_h} \\ &\leq \frac{2\epsilon}{a - \epsilon} \cdot \alpha \\ &< \beta, \end{aligned}$$

where we use the fact that $M_{i_h} \leq \alpha$ and the definition of ϵ .

Starting from (EC.1), consider $h \in \mathcal{H}$ such that $t_h \leq T_f$ and suppose by contradiction that $\pi_f^*(h) = j_h \notin \mathcal{T}(i_h)$. Even though the overall policy can be fairly complicated (history-dependent), the first action is deterministic and incurs an expected cost of $m_{i_h j_h}$ because the costs are independent across time and arcs. Moreover, when the objective is to minimize the average cost, the optimal strategy among all history-dependent rules is to follow the shortest path with respect to the mean arc costs (once again because the costs are independent across time and arcs). As a result:

$$\mathbb{E}[X_{\pi_f^*}^h] \geq m_{i_h j_h} + M_{j_h},$$

which implies:

$$\mathbb{E}[X_{\pi_f^*}^h] - M_{i_h} \geq m_{i_h j_h} + M_{j_h} - M_{i_h},$$

a contradiction. \square

B.2. Proof of Proposition 1

Proof of Proposition 1. Using Theorem 1, the optimization problem (1) can be equivalently formulated as a discrete-time finite-horizon MDP in the extended space state $(i, t) \in \mathcal{V} \times [T - \delta^{\sup}, \lceil \frac{T-T_f}{\delta^{\inf}} \rceil, T]$ where i is the current location and t is the, possibly negative, remaining budget. Specifically:

- The time horizon is $\lceil \frac{T-T_f}{\delta^{\inf}} \rceil$,
- The initial state is (s, T) ,
- The set of available actions at state (i, t) , for $i \neq d$, is taken as $\mathcal{V}(i)$. Picking $j \in \mathcal{V}(i)$ corresponds to crossing link (i, j) and results in a transition to state $(j, t - \omega)$ with probability $p_{ij}(\omega)d\omega$,
- The only available action at a state (d, t) is to remain in this state,
- The transition rewards are all equal to 0,
- The final reward at the epoch $\lceil \frac{T-T_f}{\delta^{\inf}} \rceil$ for any state (i, t) is equal to $f_i(t)$, which is the optimal expected objective-to-go when following the shortest path tree \mathcal{T} starting at node i with remaining budget t . Specifically, the collection of functions $(f_i(\cdot))_{i \in \mathcal{V}}$ is a solution to the following program:

$$\begin{aligned} f_d(t) &= f(t), & t &\leq T_f, \\ f_i(t) &= \max_{j \in \mathcal{T}(i)} \int_0^\infty p_{ij}(\omega) \cdot f_j(t - \omega) d\omega & i \neq d, t \leq T_f. \end{aligned}$$

Observe that Theorem 1 is crucial to be able to define the final rewards. Proposition 4.4.3 of Puterman (2014) shows that any *Markov* policy solution to (4) is an optimal solution to (1).

\square

B.3. Proof of Proposition 2

Proof of Proposition 2. For any node $i \in \mathcal{V}$, and $t \leq T$, we denote by $u_i^{\pi^{\Delta t}}(t)$ the expected risk function when following policy $\pi^{\Delta t}$ starting at i with remaining budget t . We deal with each case separately.

Case 1. We use the following useful facts:

- The functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are non-decreasing,
- The functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ are non-decreasing,
- The functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ lower bound the functions $(u_i(\cdot))_{i \in \mathcal{V}}$.

The main difficulty in proving convergence lies in the fact that the approximation $u_i^{\Delta t}(t)$ may not necessarily improve as Δt decreases. However, this is the case for regular mesh size sequences such as $(\Delta t_p = \frac{1}{2^p})_{p \in \mathbb{N}}$. Hence, we first demonstrate convergence in that particular case in Lemma EC.2 and rely on this last result to prove pointwise convergence in general in Lemma EC.3.

LEMMA EC.2. *For the regular mesh $(\Delta t_p = \frac{1}{2^p})_{p \in \mathbb{N}}$, the sequence $(u_i^{\Delta t_p}(t))_{p \in \mathbb{N}}$ converges to $u_i(t)$ for almost every point t in $[k_i^{\min} \cdot \Delta t, T]$.*

Proof First observe that, for any t , the sequence $(u_i^{\Delta t_p}(t))_{p \in \mathbb{N}}$ is non-decreasing since (i) the discretization mesh used at step $p+1$ is strictly contained in the discretization mesh used at step p and (ii) the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are non-decreasing. This shows that the functions $(u_i^{\Delta t_p}(\cdot))_{i \in \mathcal{V}}$ converge pointwise to some limits $(f_i(\cdot))_{i \in \mathcal{V}}$. Using the preliminary remarks, we get:

$$f_i(t) \leq u_i(t) \quad \forall t \in [k_i^{\min} \cdot \Delta t, T], \forall i \in \mathcal{V}.$$

Next, we establish that for any $i \in \mathcal{V}$, $t \in [k_i^{\min} \cdot \Delta t, T]$ and $\epsilon > 0$, $f_i(t) \geq u_i(t - \epsilon)$. This will enable us to squeeze $f_i(t)$ to finally derive $f_i(t) = u_i(t)$. We start with node d . Observe that, by construction of the approximation, $u_d^{\Delta t}(\cdot)$ converges pointwise to $f(\cdot)$ at every point of continuity of $f(\cdot)$. Furthermore, since $f_d(\cdot)$ and $u_d(\cdot)$ are non-decreasing, we have $f_d(t) \geq u_d(t - \epsilon)$ for all $t \in [k_d^{\min} \cdot \Delta t, T]$ and for all $\epsilon > 0$. Consider $\epsilon > 0$ and a large enough p such that $\epsilon > \frac{1}{2^p}$ which implies $\Delta t_p \cdot \lfloor \frac{t}{\Delta t_p} \rfloor \geq t - \epsilon$. We first show by induction on the level of the nodes in \mathcal{T} that:

$$f_i(t) \geq u_i(t - \text{level}(i, \mathcal{T}) \cdot \epsilon) \quad \forall t \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t), \forall i \in \mathcal{V}.$$

The base case follows from the discussion above. Assume that the induction property holds for all nodes of level less than l and consider a node $i \in \mathcal{V}$ of level $l+1$. We have, for $t \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t)$:

$$\begin{aligned} u_i^{\Delta t_p}(t) &= u_i^{\Delta t_p}\left(\left\lfloor \frac{t}{\Delta t_p} \right\rfloor \cdot \Delta t_p\right) \\ &\geq \max_{j \in \mathcal{T}(i)} \mathbb{E}[u_j^{\Delta t_p}\left(\left\lfloor \frac{t}{\Delta t_p} \right\rfloor \cdot \Delta t_p - c_{ij}\right)] \\ &\geq \max_{j \in \mathcal{T}(i)} \mathbb{E}[u_j^{\Delta t_p}(t - \epsilon - c_{ij})]. \end{aligned}$$

To take the limit $p \rightarrow \infty$ in the previous inequality, note that, for any $j \in \mathcal{T}(i)$:

$$u_j^{\Delta t p}(t - \epsilon - c_{ij}) \geq u_j^{\Delta t 1}(t - \epsilon - \delta^{\sup}),$$

while $(u_j^{\Delta t p}(t - \epsilon - c_{ij}))_{p \in \mathbb{N}}$ is non-decreasing and converges almost surely to $f_j(t - \epsilon - c_{ij})$ as $p \rightarrow \infty$.

Therefore, we can apply the monotone convergence theorem and derive:

$$f_i(t) \geq \mathbb{E}[f_j(t - \epsilon - c_{ij})].$$

As the last inequality holds for any $j \in \mathcal{T}(i)$, we finally obtain:

$$f_i(t) \geq \max_{j \in \mathcal{T}(i)} \mathbb{E}[f_j(t - \epsilon - c_{ij})] \quad \forall t \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t).$$

Using the induction property along with Theorem 1, we get:

$$\begin{aligned} f_i(t) &\geq \max_{j \in \mathcal{T}(i)} \mathbb{E}[u_j(t - \text{level}(i, \mathcal{T}) \cdot \epsilon - c_{ij})] \\ &\geq u_i(t - \text{level}(i, \mathcal{T}) \cdot \epsilon), \end{aligned}$$

for all $t \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t)$, which concludes the induction. We can now prove by induction on m , along the same lines as above, that:

$$f_i(t) \geq u_i(t - (|\mathcal{V}| + m) \cdot \epsilon) \quad \forall t \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t + m \cdot \delta^{\inf}), \forall i \in \mathcal{V},$$

for all $m \in \mathbb{N}$. This last result can be reformulated as:

$$f_i(t) \geq u_i(t - \epsilon) \quad \forall \epsilon > 0, \forall t \in [k_i^{\min} \cdot \Delta t, T], \forall i \in \mathcal{V}.$$

Combining this lower bound with the upper bound previously derived, we get:

$$u_i(t) \geq f_i(t) \geq u_i(t^-) \quad \forall t \in [k_i^{\min} \cdot \Delta t, T], \forall i \in \mathcal{V},$$

where $u_i(t^-)$ refers to the left one-sided limit of $u_i(\cdot)$ at t . Since, $u_i(\cdot)$ is non-decreasing, it has countably many discontinuity points and the last inequality shows that $f_i(\cdot) = u_i(\cdot)$ almost everywhere on $[k_i^{\min} \cdot \Delta t, T]$. \square

LEMMA EC.3. *For any sequence $(\Delta t_p)_{p \in \mathbb{N}}$ converging to 0, the sequence $(u_i^{\Delta t p}(t))_{p \in \mathbb{N}}$ converges to $u_i(t)$ for almost every point t in $[k_i^{\min} \cdot \Delta t, T]$.*

Proof In contrast to the particular case handled by Lemma EC.2, our approximation of $u_i(t)$ may not improve as p increases. For that reason, there is no straightforward comparison between $(u_i^{\Delta t p}(t))_p$ and $u_i(t)$. However, for a given $i \in \mathcal{V}$, $t \in [k_i^{\min} \cdot \Delta t, T]$, $\epsilon > 0$ and a large enough p , $(u_i^{\Delta t p}(t))_p$ can be shown to be lower bounded by a subsequence of $(u_i^{\frac{1}{2^p}}(t - \epsilon))_p$. This is how we proceed to establish convergence.

Consider $i \in \mathcal{V}$, $t \in [k^{\min} \cdot \Delta t, T]$, $\epsilon > 0$ and $p \in \mathbb{N}$. Define $\sigma(p) \in \mathbb{N}$ as the unique integer satisfying $\frac{1}{2^{\sigma(p)-1}} < \Delta t_p \leq \frac{1}{2^{\sigma(p)}}$. Since $\lim_{p \rightarrow \infty} \Delta t_p = 0$, we necessarily have $\lim_{p \rightarrow \infty} \sigma(p) = \infty$. Remark that $u_i^{\frac{1}{2^{\sigma(p)}}}(\cdot)$ has steps of size $\frac{1}{2^{\sigma(p)}} \geq \Delta t_p$, i.e. $u_i^{\Delta t_p}(\cdot)$ is expected to be a tighter approximation of $u_i(\cdot)$ than $u_i^{\frac{1}{2^{\sigma(p)}}}(\cdot)$ is. However, the time steps do not overlap (multiples of either Δt_p or $\frac{1}{2^p}$) making the two sequences impossible to compare. Nevertheless, the time steps differ by no more than Δt_p . Thus, if p is large enough so that $\Delta t_p < \epsilon$, for each update needed to calculate $u_i^{\frac{1}{2^{\sigma(p)}}}(t - \epsilon)$, there is a corresponding update for a larger budget to compute $u_i^{\Delta t_p}(t)$. As a consequence, the sequence $(u_i^{\frac{1}{2^{\sigma(p)}}}(t - \epsilon))_p$ constitutes a lower bound on the sequence of interest $(u_i^{\Delta t_p}(t))_p$. Using the preliminary remarks, we are able to squeeze $(u_i^{\Delta t_p}(t))_p$:

$$u_i^{\frac{1}{2^{\sigma(p)}}}(t - \epsilon) \leq u_i^{\Delta t_p}(t) \leq u_i(t),$$

for all $i \in \mathcal{V}$, $t \in [k^{\min} \cdot \Delta t, T]$, $\epsilon > 0$ and for p large enough. This can be proved first by induction on the level of the nodes in \mathcal{T} and then by interval increments of size δ^{\inf} along the same lines as what is done in Lemma EC.2. Yet, Lemma EC.2 shows that:

$$\lim_{p \rightarrow \infty} u_i^{\frac{1}{2^{\sigma(p)}}}(t - \epsilon) = u_i(t - \epsilon),$$

provided $t - \epsilon$ is a point of continuity for $u_i(\cdot)$. As $u_i(\cdot)$ has countably many discontinuity points (it is non-decreasing), the last inequality shows, by taking p large enough and ϵ small enough, that $u_i^{\Delta t_p}(t) \rightarrow_{p \rightarrow \infty} u_i(t)$ for t a point of continuity of $u_i(\cdot)$. \square

Case 2. The first step consists in proving that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T]$. By induction on l , we start by proving that $u_i(\cdot)$ is continuous on $(-\infty, T_f)$ for all nodes i of level l in \mathcal{T} . The base case follows from the continuity of $f(\cdot)$. Assuming the property holds for some $l \geq 1$, we consider a node i of level $l+1$ in \mathcal{T} , $t < T_f$ and a sequence $t_n \rightarrow_{n \rightarrow \infty} t$. Using Theorem 1, we have:

$$|u_i(t) - u_i(t_n)| \leq \max_{j \in \mathcal{T}(i)} \mathbb{E}[|u_j(t - c_{ij}) - u_j(t_n - c_{ij})|].$$

For any $j \in \mathcal{T}(i)$, we can use the uniform continuity of $u_j(\cdot)$ on $[t - 2 \cdot \delta^{\sup}, t]$ to prove that this last term converges to 0 as $n \rightarrow \infty$. We conclude that all the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T_f)$. By induction on m , we can then show that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T_f + m \cdot \delta^{\inf})$, to finally conclude that they are continuous on $(-\infty, T]$. We are now able to prove uniform convergence. Since $[T_f - |\mathcal{V}| \cdot \delta^{\sup}, T]$ is a compact set, the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are also uniformly continuous on this set. Take $\epsilon > 0$, there exists $\alpha > 0$ such that:

$$\forall i \in \mathcal{V}, |u_i(\omega) - u_i(\omega')| \leq \epsilon, \forall (\omega, \omega') \in [T_f - |\mathcal{V}| \cdot \delta^{\sup}, T]^2 \text{ with } |\omega - \omega'| \leq \alpha.$$

Building on this, we can show, by induction on the level of the nodes in \mathcal{T} , that:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t)} |u_i(\omega) - u_i^{\Delta t}(\omega)| \leq \text{level}(i, \mathcal{T}) \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

This follows from the sequence of inequalities:

$$\begin{aligned} \sup_{\omega \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t)} |u_i^{\Delta t}(\omega) - u_i(\omega)| &\leq \sup_{k \in \{k_i^{\min}, \dots, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor - 1\}} |u_i^{\Delta t}(k \cdot \Delta t) - u_i(k \cdot \Delta t)| \\ &+ \sup_{\omega \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i(\omega) - u_i(\left\lfloor \frac{\omega}{\Delta t} \right\rfloor \cdot \Delta t)| \\ &\leq \sup_{k \in \{k_i^{\min}, \dots, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor - 1\}} \max_{j \in \mathcal{T}(i)} \mathbb{E}[|u_j^{\Delta t}(k \cdot \Delta t - c_{ij}) - u_j(k \cdot \Delta t - c_{ij})|] \\ &+ \epsilon \\ &\leq (\text{level}(i, \mathcal{T}) - 1) \cdot \epsilon + \epsilon \\ &\leq \text{level}(i, \mathcal{T}) \cdot \epsilon. \end{aligned}$$

We conclude that:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t)} |u_i(\omega) - u_i^{\Delta t}(\omega)| \leq |\mathcal{V}| \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

Along the same lines, we can show by induction on m that:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t + m \cdot \delta^{\inf})} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq (|\mathcal{V}| + m) \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

This implies:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, T]} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\inf}} \right\rceil + 1) \cdot \epsilon, \quad \forall i \in \mathcal{V},$$

assuming $\Delta t \leq \delta^{\inf}$. In particular, this shows uniform convergence. To conclude the proof of Case 2, we show that $\pi^{\Delta t}$ is a $o(1)$ -approximate optimal solution to (1) as $\Delta t \rightarrow 0$. Using the last set of inequalities derived in combination with the uniform continuity of the functions $(u_i(\cdot))_{i \in \mathcal{V}}$, we can show that:

$$\forall i \in \mathcal{V}, |u_i^{\Delta t}(\omega) - u_i^{\Delta t}(\omega')| \leq (2 \cdot |\mathcal{V}| + 2 \cdot \left\lceil \frac{T - T_f}{\delta^{\inf}} \right\rceil + 3) \cdot \epsilon, \quad \forall (\omega, \omega') \in [k_i^{\min} \cdot \Delta t, T]^2 \text{ with } |\omega - \omega'| \leq \alpha, \quad (\text{EC.2})$$

and:

$$\forall i \in \mathcal{V}, |u_i^{\Delta t}(\omega) - u_i(\omega')| \leq (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\inf}} \right\rceil + 2) \cdot \epsilon, \quad \forall (\omega, \omega') \in [k_i^{\min} \cdot \Delta t, T]^2 \text{ with } |\omega - \omega'| \leq \alpha. \quad (\text{EC.3})$$

We can now prove, by induction on the level of the nodes in \mathcal{T} , that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 3 \cdot \text{level}(i, \mathcal{T}) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon, \quad \forall t \in [k_i^{\text{min}} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t), \forall i \in \mathcal{V}.$$

This follows from the sequence of inequalities:

$$\begin{aligned} u_i^{\pi^{\Delta t}}(t) &= \int_0^\infty p_{i\pi^{\Delta t}(i,t)}(\omega) \cdot u_{\pi^{\Delta t}(i,t)}^{\pi^{\Delta t}}(t - \omega) d\omega \\ &\geq \int_0^\infty p_{i\pi^{\Delta t}(i,t)}(\omega) \cdot u_{\pi^{\Delta t}(i,t)}(t - \omega) d\omega - 3 \cdot (\text{level}(i, \mathcal{T}) - 1) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon \\ &\geq \int_0^\infty p_{i\pi^{\Delta t}(i,t)}(\omega) \cdot u_{\pi^{\Delta t}(i,t)}^{\Delta t}(t - \omega) d\omega - \epsilon - 3 \cdot (\text{level}(i, \mathcal{T}) - 1) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon \\ &\geq \int_0^\infty p_{i\pi^{\Delta t}(i,t)}(\omega) \cdot u_{\pi^{\Delta t}(i, \lfloor \frac{t}{\Delta t} \rfloor \cdot \Delta t)}^{\Delta t} \left(\left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t - \omega \right) d\omega - \epsilon - (2 \cdot |\mathcal{V}| + 2 \cdot \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 3) \cdot \epsilon \\ &\quad - 3 \cdot (\text{level}(i, \mathcal{T}) - 1) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon \\ &\geq u_i^{\Delta t} \left(\left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t \right) - 3 \cdot \text{level}(i, \mathcal{T}) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon, \end{aligned}$$

where we use the induction property for the first inequality, the uniform convergence for the second, (EC.2) for the third, the definition of $\pi^{\Delta t}(i, t)$ for the fourth and finally (EC.3). We conclude that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 3 \cdot |\mathcal{V}| \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon, \quad \forall t \in [k_i^{\text{min}} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t), \forall i \in \mathcal{V}.$$

We can then prove by induction on m , in the same fashion as above, that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 3 \cdot (m + |\mathcal{V}|) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2) \cdot \epsilon, \quad \forall t \in [k_i^{\text{min}} \cdot \Delta t, \left\lfloor \frac{T_f}{\Delta t} \right\rfloor \cdot \Delta t + m \cdot \delta^{\text{inf}}), \forall i \in \mathcal{V},$$

for all m . We conclude that $u_s^{\pi^{\Delta t}}(T) \geq u_s(T) - 3 \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2)^2 \cdot \epsilon$, which establishes the claim.

Case 3. The first step consists in showing that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are Lipschitz on $[T_f - |\mathcal{V}| \cdot \delta^{\text{sup}}, T]$. Take K to be a Lipschitz constant for $f(\cdot)$ on $[T_f - (\left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2 \cdot |\mathcal{V}| - 1) \cdot \delta^{\text{sup}}, T]$. We first show by induction on l that $u_i(\cdot)$ is K -Lipschitz on $[T_f - (\left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2 \cdot |\mathcal{V}| - l) \cdot \delta^{\text{sup}}, T_f]$ for all nodes i of level l in \mathcal{T} . The base case follows from the definition of K . Assuming the property holds for some $l \geq 1$, we consider a node i of level $l + 1$ in \mathcal{T} . Using Theorem 1, we have, for $(t, t') \in [T_f - (\left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + 2 \cdot |\mathcal{V}| - l - 1) \cdot \delta^{\text{sup}}, T_f]^2$:

$$\begin{aligned} |u_i(t) - u_i(t')| &\leq \max_{j \in \mathcal{T}(i)} \mathbb{E}[|u_j(t - \omega) - u_j(t' - \omega)|] \\ &\leq K \cdot |t - t'|, \end{aligned}$$

where we use the induction property for l (recall that $p_{ij}(\omega) = 0$ for $\omega \geq \delta^{\text{sup}}$). We conclude that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are all K -Lipschitz on $[T_f - (\left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil + |\mathcal{V}|) \cdot \delta^{\text{sup}}, T_f]$. We now prove, by induction on m , that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are all K -Lipschitz on $[T_f - (\left\lceil \frac{T - T_f}{\delta^{\text{inf}}} \right\rceil - m + |\mathcal{V}|) \cdot \delta^{\text{sup}}, T_f + m \cdot \delta^{\text{inf}}]$.

The base case follows from the previous induction. Assuming the property holds for some m , we have for $i \in \mathcal{V}$ and for $(t, t') \in [T_f - (\lceil \frac{T-T_f}{\delta^{\inf}} \rceil - m - 1 + |\mathcal{V}|) \cdot \delta^{\sup}, T_f + (m+1) \cdot \delta^{\inf}]^2$:

$$\begin{aligned} |u_i(t) - u_i(t')| &\leq \max_{j \in \mathcal{V}(i)} \mathbb{E}[|u_j(t - \omega) - u_j(t' - \omega)|] \\ &\leq K \cdot |t - t'|, \end{aligned}$$

where we use the fact that $p_{ij}(\omega) = 0$ for $\omega \leq \delta^{\inf}$ or $\omega \geq \delta^{\sup}$ and the induction property. We conclude that the function $(u_i(\cdot))_{i \in \mathcal{V}}$ are all K -Lipschitz on $[T_f - |\mathcal{V}| \cdot \delta^{\sup}, T]$. Using this last fact, we can prove, by induction on the level of the nodes in \mathcal{T} , in a similar fashion as done for Case 2, that:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, \lfloor \frac{T_f}{\Delta t} \rfloor \cdot \Delta t)} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq \text{level}(i, \mathcal{T}) \cdot K \cdot \Delta t, \quad \forall i \in \mathcal{V}.$$

By induction on m , we can then show that:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, \lfloor \frac{T_f}{\Delta t} \rfloor \cdot \Delta t + m \cdot \delta^{\inf})} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq (|\mathcal{V}| + m) \cdot K \cdot \Delta t, \quad \forall i \in \mathcal{V}, \forall m \in \mathbb{N}.$$

This implies:

$$\sup_{\omega \in [k_i^{\min} \cdot \Delta t, T]} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq (|\mathcal{V}| + \lceil \frac{T-T_f}{\delta^{\inf}} \rceil + 1) \cdot K \cdot \Delta t, \quad \forall i \in \mathcal{V},$$

assuming $\Delta t \leq \delta^{\inf}$. This shows uniform convergence at speed Δt . To conclude the proof of Case 3, we show that $\pi^{\Delta t}$ is a $O(\Delta t)$ -approximate optimal solution to (1) as $\Delta t \rightarrow 0$. We can show, using the last inequality derived along with the same sequence of inequalities as in Case 2, by induction on the level of the nodes in \mathcal{T} that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 6 \cdot \text{level}(i, \mathcal{T}) \cdot (|\mathcal{V}| + \lceil \frac{T-T_f}{\delta^{\inf}} \rceil + 1) \cdot K \cdot \Delta t, \quad \forall t \in [k_i^{\min} \cdot \Delta t, \lfloor \frac{T_f}{\Delta t} \rfloor \cdot \Delta t), \forall i \in \mathcal{V}.$$

We can then prove by induction on m , in the same fashion as in Case 2, that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 6 \cdot (m + |\mathcal{V}|) \cdot (|\mathcal{V}| + \lceil \frac{T-T_f}{\delta^{\inf}} \rceil + 1) \cdot K \cdot \Delta t, \quad \forall t \in [k_i^{\min} \cdot \Delta t, \lfloor \frac{T_f}{\Delta t} \rfloor \cdot \Delta t + m \cdot \delta^{\inf}), \forall i \in \mathcal{V},$$

for all $m \in \mathbb{N}$. We conclude that $u_s^{\pi^{\Delta t}}(T) \geq u_s(T) - 6 \cdot (|\mathcal{V}| + \lceil \frac{T-T_f}{\delta^{\inf}} \rceil + 1)^2 \cdot K \cdot \Delta t$, which establishes the claim.

Case 4. In this situation, we can show by induction on m that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $[0, m \cdot \delta^{\inf}]$ and conclude that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $[0, T]$. Since $[0, T]$ is a compact set, these functions are also uniformly continuous on $[0, T]$. Moreover, $u_i^{\Delta t}(t) = u_i(t) = 0$ for all $t \leq 0$ and $i \in \mathcal{V}$. Using these two observations, we can apply the same techniques as in Case 2 to obtain the same results. \square

B.4. Proof of Theorem 2

Proof of Theorem 2. As in Theorem 1, the last part of the theorem is trivial because any strategy is optimal when having already spent a budget of $T - T_f^r$.

The proof for the first part is an extension of Theorem 1 and follows the same steps. We denote by $(m_{ij})_{(i,j) \in \mathcal{A}}$ the worst-case expected costs, i.e.:

$$m_{ij} = \sup_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[X] \quad \forall (i, j) \in \mathcal{A}.$$

Observe that these quantities are well-defined as \mathcal{P}_{ij} is not empty and $\mathbb{E}_{X \sim p_{ij}}[X] \leq \delta^{\sup}$ for any $p_{ij} \in \mathcal{P}_{ij}$. Furthermore, there exists $p_{ij}^* \in \mathcal{P}_{ij}$ such that $m_{ij} = \mathbb{E}_{X \sim p_{ij}^*}[X]$ as \mathcal{P}_{ij} is compact for the weak topology. For any node $i \neq d$, we define M_i as the length of a shortest path from i to d in \mathcal{G} when the arc costs are taken as $(m_{ij})_{(i,j) \in \mathcal{A}}$. Just like in Theorem 1, we consider an optimal strategy $\pi_{f,\mathcal{P}}^*$ solution to (3). For a given history $h \in \mathcal{H}$, we define t_h as the remaining budget, i.e. T minus the total cost spent so far, and i_h as the current location. The policy $\pi_{f,\mathcal{P}}^*$ maps $h \in \mathcal{H}$ to a probability distribution over $\mathcal{V}(i_h)$. Observe that randomizing does not help because (i) the costs are independent across time and arcs and (ii) the ambiguity set is rectangular. Hence, without loss of generality, we may assume that $\pi_{f,\mathcal{P}}^*$ actually maps h to the node in $\mathcal{V}(i_h)$ minimizing the worst-case objective function given h . For $h \in \mathcal{H}$, we denote by $X_{\pi_{f,\mathcal{P}}^*}^h$ the random cost-to-go incurred by following strategy $\pi_{f,\mathcal{P}}^*$, i.e. not including the total cost spent up to this point of the history $T - t_h$. We define π_s as a policy associated with an arbitrary shortest path from i to d with respect to $(m_{ij})_{(i,j) \in \mathcal{A}}$. Specifically, π_s maps the current location i_h to a node in $\mathcal{T}^r(i_h)$, irrespective of the history of the process. Similarly as for $\pi_{f,\mathcal{P}}^*$, we denote by $X_{\pi_s}^h$ the random cost-to-go incurred by following strategy π_s for $h \in \mathcal{H}$. Using *Bellman's Principle of Optimality* for $\pi_{f,\mathcal{P}}^*$, we have:

$$\begin{aligned} \mathbb{E}_{\mathbf{P}^*}[f(t_h - X_{\pi_{f,\mathcal{P}}^*}^h)] &\geq \inf_{\forall \tau \geq T - t_h, \forall (i,j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau}[f(t - X_{\pi_{f,\mathcal{P}}^*}^h)] \\ &\geq \sup_{\pi \in \Pi} \inf_{\forall \tau \geq T - t_h, \forall (i,j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau}[f(t - X_\pi^h)] \\ &\geq \inf_{\forall \tau \geq T - t_h, \forall (i,j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau}[f(t - X_{\pi_s}^h)] \\ &\geq \mathbb{E}_{\mathbf{Q}^\tau}[f(t - X_{\pi_s}^h)], \end{aligned}$$

where $(q_{ij}^\tau)_{(i,j) \in \mathcal{A}, \tau \geq T - t_h}$ is given by the worst-case scenario in the ambiguity sets, i.e.:

$$(q_{ij}^\tau)_{(i,j) \in \mathcal{A}, \tau \geq T - t_h} \in \arg \min_{\forall \tau \geq T - t_h, \forall (i,j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau}[f(t - X_{\pi_s}^h)],$$

which can be shown to exist because the ambiguity sets are compact. Using the last inequality derived, we can prove, using the exact same sequence of inequalities as in Theorem 1, that there exists T_f^r such that, for both cases (a) and (b):

$$\mathbb{E}_{\mathbf{P}^*}[X_{\pi_{f,\mathcal{P}}^*}^h] - \mathbb{E}_{\mathbf{Q}^\tau}[X_{\pi_s}^h] < \min_{i \neq d} \min_{j \in \mathcal{V}(i), j \notin \mathcal{T}^r(i)} \{m_{ij} + M_j - M_i\} \quad \forall h \in \mathcal{H} \text{ such that } t_h \leq T_f^r, \quad (\text{EC.4})$$

with the convention that the minimum of an empty set is equal to infinity. Starting from (EC.4), consider $h \in \mathcal{H}$ such that $t_h \leq T_f^r$ and suppose by contradiction that $\pi_{f,\mathcal{P}}^*(h) = j_h \notin \mathcal{T}^r(i_h)$. As mentioned in Theorem 1, even though $\pi_{f,\mathcal{P}}^*$ can be fairly complicated, the first action is deterministic and incurs an expected cost of $m_{i_h j_h}$ because the costs are independent across time and arcs. Moreover, when the objective is to minimize the average cost, the optimal strategy among all history-dependent rules is to follow the shortest path with respect to the mean arc costs (once again because the costs are independent across time and arcs). As a result:

$$\mathbb{E}_{\mathbf{p}^*}[X_{\pi_{f,\mathcal{P}}^*}^h] \geq m_{i_h j_h} + M_{j_h}.$$

Additionally, by definition of $(p_{ij}^*)_{(i,j) \in \mathcal{A}}$:

$$\begin{aligned} \mathbb{E}_{\mathbf{q}^\tau}[X_{\pi_s}^h] &\leq \mathbb{E}_{\mathbf{p}^*}[X_{\pi_s}^h] \\ &\leq M_{i_h}. \end{aligned}$$

This implies:

$$\mathbb{E}_{\mathbf{p}^*}[X_{\pi_{f,\mathcal{P}}^*}^h] - \mathbb{E}_{\mathbf{q}^\tau}[X_{\pi_s}^h] \geq m_{i_h j_h} + M_{j_h} - M_{i_h},$$

a contradiction. We conclude that:

$$\pi_{f,\mathcal{P}}^*(h) \in \mathcal{T}^r(i_h) \quad \forall h \in \mathcal{H} \text{ such that } t_h \leq T_f^r.$$

□

B.5. Proof of Proposition 3.

Proof of Proposition 3. The proof uses a reduction to distributionally robust finite-horizon MDPs in a similar fashion as in Proposition 1. Using Theorem 2, the optimization problem (3) can be equivalently formulated as a discrete-time finite-horizon distributionally robust MDP in the extended space state $(i, t) \in \mathcal{V} \times [T - \delta^{\sup} \cdot \left\lceil \frac{T - T_f^r}{\delta^{\inf}} \right\rceil, T]$ where i is the current location and t is the, possibly negative, remaining budget. Specifically:

- The time horizon is $\left\lceil \frac{T - T_f^r}{\delta^{\inf}} \right\rceil$,
- The initial state is (s, T) ,
- The set of available actions at state (i, t) , for $i \neq d$, is taken as $\mathcal{V}(i)$. Picking $j \in \mathcal{V}(i)$ corresponds to crossing link (i, j) and results in a transition to state $(j, t - \omega)$ with probability $p_{ij}(\omega)d\omega$,
- The probability of transitions are only known to lie in the rectangular ambiguity set:

$$\prod_{\substack{(i,j) \in \mathcal{A} \\ t \in [T - \delta^{\sup} \cdot \left\lceil \frac{T - T_f^r}{\delta^{\inf}} \right\rceil, T]}} \mathcal{P}_{ij},$$

- The only available action at a state (d, t) is to remain in this state,
- The transition rewards are all equal to 0,
- The final reward at the epoch $\left\lceil \frac{T - T_f^r}{\delta_{\inf}} \right\rceil$ for any state (i, t) is equal to $f_i(t)$, which is the optimal worst-case expected objective-to-go when following the shortest path tree \mathcal{T}^r starting at node i with remaining budget t . Specifically, the collection of functions $(f_i(\cdot))_{i \in \mathcal{V}}$ is a solution to the following program:

$$\begin{aligned} f_d(t) &= f(t), & t &\leq T_f^r \\ f_i(t) &= \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot f_j(t - \omega) d\omega & i \neq d, t \leq T_f^r. \end{aligned}$$

As a consequence, we can conclude the proof with Theorem 2.2 of Iyengar (2005) (or equivalently Theorem 1 of Nilim and Ghaoui (2005)). \square

B.6. Proof of Proposition 4

The proofs are along the same lines as for Proposition 2.

Proof of Proposition 4. We deal with each case separately.

Case 1. We make use the following facts:

- The functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ are non-decreasing,
- The functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are non-decreasing,
- The functions $(u_i^{\Delta t}(\cdot))_{i \in \mathcal{V}}$ lower bound the functions $(u_i(\cdot))_{i \in \mathcal{V}}$.

We follow the same recipe as in Proposition 2. We start by proving convergence for the discretization sequence $(\Delta t_p = \frac{1}{2^p})_{p \in \mathbb{N}}$. Then, we conclude the general study with the exact same argument as in Lemma EC.3.

LEMMA EC.4. *For the regular mesh $(\Delta t_p = \frac{1}{2^p})_{p \in \mathbb{N}}$, the sequence $(u_i^{\Delta t_p}(t))_{p \in \mathbb{N}}$ converges to $u_i(t)$ for almost every point t in $[k_i^{r, \min} \cdot \Delta t, T]$.*

Proof Just like in Lemma EC.2 we can prove that the sequence $(u_i^{\Delta t_p}(t))_{p \in \mathbb{N}}$ is non-decreasing for any t and $i \in \mathcal{V}$. Hence, the functions $(u_i^{\Delta t_p}(\cdot))_{i \in \mathcal{V}}$ converge pointwise to some limits $(f_i(\cdot))_{i \in \mathcal{V}}$. Using the preliminary remarks, we get:

$$f_i(t) \leq u_i(t) \quad \forall t \in [k_i^{r, \min} \cdot \Delta t, T], \forall i \in \mathcal{V}.$$

Next, we establish that for any $t \in [k_i^{r, \min} \cdot \Delta t, T]$ and for any $\epsilon > 0$, $f_i(t) \geq u_i(t - \epsilon)$. This will enable us to squeeze $f_i(t)$ to finally derive $f_i(t) = u_i(t)$. We start with node d . Observe that, by construction of the approximation, $u_d^{\Delta t}(\cdot)$ converges pointwise to $f(\cdot)$ at every point of continuity of $f(\cdot)$. Furthermore, since $f_d(\cdot)$ and $u_d(\cdot)$ are non-decreasing, we have $f_d(t) \geq u_d(t - \epsilon)$ for all $t \in [k_d^{r, \min} \cdot \Delta t, T]$ and for

all $\epsilon > 0$. Consider $\epsilon > 0$ and a large enough p such that $\epsilon > \frac{1}{2^p}$ which implies $\Delta t_p \cdot \lfloor \frac{t}{\Delta t_p} \rfloor \geq t - \epsilon$. We first show by induction on the level of the nodes in \mathcal{T}^r that:

$$f_i(t) \geq u_i(t - \text{level}(i, \mathcal{T}^r) \cdot \epsilon) \quad \forall t \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t], \forall i \in \mathcal{V}.$$

The base case follows from the discussion above. Assume that the induction property holds for all nodes of level less than l and consider a node $i \in \mathcal{V}$ of level $l+1$. We have, for $t \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]$:

$$\begin{aligned} u_i^{\Delta t_p}(t) &\geq u_i^{\Delta t_p}\left(\left\lfloor \frac{t}{\Delta t_p} \right\rfloor \cdot \Delta t_p\right) \\ &\geq \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [u_j^{\Delta t_p}\left(\left\lfloor \frac{t}{\Delta t_p} \right\rfloor \cdot \Delta t_p - X\right)] \\ &\geq \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [u_j^{\Delta t_p}(t - \epsilon - X)]. \end{aligned}$$

Take $j \in \mathcal{T}^r(i)$. Since $u_j^{\Delta t_p}(\cdot)$ is continuous and \mathcal{P}_{ij} is compact, the infimum in the previous inequality is attained for some $p_{ij}^p \in \mathcal{P}_{ij}$ which gives:

$$u_i^{\Delta t_p}(t) \geq \mathbb{E}_{X \sim p_{ij}^p} [u_j^{\Delta t_p}(t - \epsilon - X)].$$

As the sequence $(u_j^{\Delta t_p}(t - \epsilon - \omega))_p$ is non-decreasing for any ω , we have, for any $m \leq p$:

$$u_i^{\Delta t_p}(t) \geq \mathbb{E}_{X \sim p_{ij}^p} [u_j^{\Delta t_m}(t - \epsilon - X)].$$

Because \mathcal{P}_{ij} is a compact set for the weak topology, there exists a subsequence of $(p_{ij}^p)_p$ converging weakly in \mathcal{P}_{ij} to some probability measure p_{ij} . Without loss of generality, we continue to refer to this subsequence as $(p_{ij}^p)_p$. We can now take the limit $p \rightarrow \infty$ in the previous inequality which yields:

$$f_i(t) \geq \mathbb{E}_{X \sim p_{ij}} [u_j^{\Delta t_m}(t - \epsilon - X)],$$

since $u_j^{\Delta t_m}(\cdot)$ is continuous. To take the limit $m \rightarrow \infty$, note that:

$$u_j^{\Delta t_m}(t - \epsilon - X) \geq u_j^{\Delta t_1}(t - \epsilon - \delta^{\sup}),$$

while $(u_j^{\Delta t_m}(t - \epsilon - X))_{m \in \mathbb{N}}$ is non-decreasing and converges almost surely to $f_j(t - \epsilon - X)$ as $m \rightarrow \infty$.

Therefore, we can apply the monotone convergence theorem and derive:

$$f_i(t) \geq \mathbb{E}_{X \sim p_{ij}} [f_j(t - \epsilon - X)],$$

which further implies

$$f_i(t) \geq \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [f_j(t - \epsilon - X)].$$

As the last inequality holds for any $j \in \mathcal{T}^r(i)$, we finally obtain:

$$f_i(t) \geq \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [f_j(t - \epsilon - X)] \quad \forall t \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t], \forall i \in \mathcal{V}.$$

Using the induction property along with Theorem 2, we get:

$$\begin{aligned} f_i(t) &\geq \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [u_j(t - \text{level}(i, \mathcal{T}^r) \cdot \epsilon - X)] \\ &\geq u_i(t - \text{level}(i, \mathcal{T}^r) \cdot \epsilon), \end{aligned}$$

for all $t \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]$, which concludes the induction. We can now prove by induction on m , along the same lines as above, that:

$$f_i(t) \geq u_i(t - (|\mathcal{V}| + m) \cdot \epsilon) \quad \forall t \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t + m \cdot \delta^{\inf}], \forall i \in \mathcal{V},$$

for all $m \in \mathbb{N}$. This last result can be reformulated as:

$$f_i(t) \geq u_i(t - \epsilon) \quad \forall \epsilon > 0, \forall t \in [k_i^{r, \min} \cdot \Delta t, T], \forall i \in \mathcal{V}.$$

Combining this lower bound with the upper bound previously derived, we get:

$$u_i(t) \geq f_i(t) \geq u_i(t^-) \quad \forall t \in [k_i^{r, \min} \cdot \Delta t, T], \forall i \in \mathcal{V},$$

where $u_i(t^-)$ refers to the left one-sided limit of $u_i(\cdot)$ at t . Since, $u_i(\cdot)$ is non-decreasing, it has countably many discontinuity points and the last inequality shows that $f_i(\cdot) = u_i(\cdot)$ almost everywhere on $[k_i^{r, \min} \cdot \Delta t, T]$. \square

Case 2. The first step consists in proving that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T]$. By induction on l , we start by proving that $u_i(\cdot)$ is continuous on $(-\infty, T_f^r]$ for all nodes i of level l in \mathcal{T}^r . The base case follows from the continuity of $f(\cdot)$. Assuming the property holds for some $l \geq 1$, we consider a node i of level $l + 1$ in \mathcal{T}^r , $t \leq T_f^r$ and a sequence $t_n \rightarrow_{n \rightarrow \infty} t$. Using Theorem 2, we have:

$$|u_i(t) - u_i(t_n)| \leq \max_{j \in \mathcal{T}^r(i)} \sup_{p_{ij} \in \mathcal{P}_{ij}} \int_0^\infty p_{ij}(\omega) \cdot |u_j(t - \omega) - u_j(t_n - \omega)| d\omega.$$

For any $j \in \mathcal{T}^r(i)$, we can use the uniform continuity of $u_j(\cdot)$ on $[t - 2 \cdot \delta^{\sup}, t]$ to prove that this last term converges to 0 as $n \rightarrow \infty$. We conclude that all the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T_f^r]$. By induction on m , we can then show that the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are continuous on $(-\infty, T_f^r + m \cdot \delta^{\inf}]$, to finally conclude that they are continuous on $(-\infty, T]$. We are now able to prove uniform convergence. Since $[T_f^r - |\mathcal{V}| \cdot \delta^{\sup}, T]$ is a compact set, the functions $(u_i(\cdot))_{i \in \mathcal{V}}$ are also uniformly continuous on this set. Take $\epsilon > 0$, there exists $\alpha > 0$ such that:

$$\forall i \in \mathcal{V}, |u_i(\omega) - u_i(\omega')| \leq \epsilon, \quad \forall (\omega, \omega') \in [T_f^r - |\mathcal{V}| \cdot \delta^{\sup}, T]^2 \text{ with } |\omega - \omega'| \leq \alpha.$$

Building on this, we can show, by induction on the level of the nodes in \mathcal{T}^r , that:

$$\sup_{\omega \in [k_i^{r, \min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i(\omega) - u_i^{\Delta t}(\omega)| \leq 2 \cdot \text{level}(i, \mathcal{T}^r) \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

This follows from the sequence of inequalities:

$$\begin{aligned}
\sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i^{\Delta t}(\omega) - u_i(\omega)| &\leq \sup_{k \in \{k_i^{r,\min}, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor\}} |u_i^{\Delta t}(k \cdot \Delta t) - u_i(k \cdot \Delta t)| \\
&+ \sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i(\omega) - u_i(\left\lfloor \frac{\omega}{\Delta t} \right\rfloor \cdot \Delta t)| \\
&+ \sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i(\omega) - u_i(\left\lceil \frac{\omega}{\Delta t} \right\rceil \cdot \Delta t)| \\
&\leq \sup_{k \in \{k_i^{r,\min}, \dots, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor\}} \max_{j \in \mathcal{T}(i)} \sup_{p_{ij} \in \mathcal{P}_{ij}} \\
&\{ \int_0^\infty p_{ij}(\omega) \cdot |u_j^{\Delta t}(k \cdot \Delta t - \omega) - u_j(k \cdot \Delta t - \omega)| d\omega \} \\
&+ 2 \cdot \epsilon \\
&\leq 2 \cdot (\text{level}(i, \mathcal{T}^r) - 1) \cdot \epsilon + 2 \cdot \epsilon \\
&\leq 2 \cdot \text{level}(i, \mathcal{T}^r) \cdot \epsilon.
\end{aligned}$$

We conclude that:

$$\sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t]} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq 2 \cdot |\mathcal{V}| \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

Along the same lines, we can show by induction on m that:

$$\sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, \left\lfloor \frac{T_f^r}{\Delta t} \right\rfloor \cdot \Delta t + m \cdot \delta^{\text{inf}}]} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq 2 \cdot (|\mathcal{V}| + m) \cdot \epsilon, \quad \forall i \in \mathcal{V}.$$

This implies:

$$\sup_{\omega \in [k_i^{r,\min} \cdot \Delta t, T]} |u_i^{\Delta t}(\omega) - u_i(\omega)| \leq 2 \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f^r}{\delta^{\text{inf}}} \right\rceil) \cdot \epsilon, \quad \forall i \in \mathcal{V},$$

assuming $\Delta t \leq \delta^{\text{inf}}$. In particular, this shows uniform convergence. To conclude the proof of Case 2, we show that $\pi^{\Delta t}$ is a $o(1)$ -approximate optimal solution to (3) as $\Delta t \rightarrow 0$. We denote by $u_i^{\pi^{\Delta t}}(t)$ the worst-case expected risk function when following policy $\pi^{\Delta t}$ starting at i with remaining budget t . We can show, by induction on the level of the nodes in \mathcal{T}^r , that :

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 12 \cdot \text{level}(i, \mathcal{T}^r) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f^r}{\delta^{\text{inf}}} \right\rceil) \cdot \epsilon, \quad \forall t \in [k_i^{r,\min} \cdot \Delta t, T_f^r], \forall i \in \mathcal{V}.$$

To do so, we can use the same sequence of inequalities as in Case 2 of Proposition 2, except that we also take the infimum over $p_{i\pi^{\Delta t}(i,t)} \in \mathcal{P}_{i\pi^{\Delta t}(i,t)}$. We derive:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 12 \cdot |\mathcal{V}| \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f^r}{\delta^{\text{inf}}} \right\rceil) \cdot \epsilon, \quad \forall t \in [k_i^{r,\min} \cdot \Delta t, T_f^r], \forall i \in \mathcal{V}.$$

Along the same lines, we can show by induction on m that:

$$u_i^{\pi^{\Delta t}}(t) \geq u_i(t) - 12 \cdot (|\mathcal{V}| + m) \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f^r}{\delta^{\inf}} \right\rceil) \cdot \epsilon, \forall t \in [k_i^{r, \min} \cdot \Delta t, T_f^r + m \cdot \delta^{\inf}], \forall i \in \mathcal{V}.$$

We conclude that $u_s^{\pi^{\Delta t}}(T) \geq u_s(T) - 12 \cdot (|\mathcal{V}| + \left\lceil \frac{T - T_f^r}{\delta^{\inf}} \right\rceil)^2 \cdot \epsilon$, which establishes the claim.

Case 3. This case is essentially identical to Case 2 substituting uniform continuity for Lipschitz continuity and the proof mirrors the proof of Case 3 of Proposition 2.

□

B.7. Proof of Lemma 1

Proof of Lemma 1. For a real value x , δ_x refers to the Dirac distribution at x . We denote by $(m_{ij})_{(i,j) \in \mathcal{A}}$ the worst-case expected costs, i.e.:

$$m_{ij} = \sup_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[X] \quad \forall (i, j) \in \mathcal{A}.$$

We define M_i as the length of a shortest path from i to d in \mathcal{G} when the arc costs are taken as $(m_{ij})_{(i,j) \in \mathcal{A}}$. Observe that $f(\cdot)$ is increasing since $f(\cdot)$ is convex and $f' \rightarrow_{-\infty} a > 0$. We use Proposition 3 and consider a solution $(\pi_{f, \mathcal{P}}^*, (u_i(\cdot))_{i \in \mathcal{V}})$ to the dynamic programming equation (7). We first prove by induction on the level of the nodes in \mathcal{T}^r that:

$$u_i(t) = f(t - M_i) \quad \forall t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\sup}, T_f^r],$$

for all nodes $i \in \mathcal{V}$. The base case is trivial. Assume that the property holds for all nodes of level less than l and consider a node $i \in \mathcal{V}$ of level $l + 1$. Take $t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\sup}, T_f^r]$. Using Theorem 2, we have:

$$\begin{aligned} u_i(t) &= \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[u_j(t - X)] \\ &= \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}}[f(t - X - M_j)] \\ &\geq \max_{j \in \mathcal{T}^r(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} f(t - \mathbb{E}_{X \sim p_{ij}}[X] - M_j) \\ &\geq \max_{j \in \mathcal{T}^r(i)} f(t - m_{ij} - M_j) \\ &\geq f(t - \min_{j \in \mathcal{T}^r(i)} m_{ij} - M_j) \\ &\geq f(t - M_i) \\ &\geq \max_{j \in \mathcal{T}^r(i)} \mathbb{E}_{X \sim \delta_{m_{ij}}} [f(t - X - M_j)], \end{aligned}$$

where the first inequality results from the convexity of $f(\cdot)$ and the third inequality is a consequence of the monotonicity of $f(\cdot)$. Since $\delta_{m_{ij}} \in \mathcal{P}_{ij}$, the last inequality shows that $u_i(t) = f(t - M_i)$. This concludes the induction. We move on to prove by induction on m that:

$$u_i(t) = f(t - M_i) \quad \forall t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\sup}, T_f^r + m \cdot \delta^{\inf}], \forall i \in \mathcal{V}.$$

Assume that the inductive property holds for some $m \in \mathbb{N}$. Consider $i \neq d$. We have, for $t \in [T_f^r + m \cdot \delta^{\text{inf}}, T_f^r + (m+1) \cdot \delta^{\text{inf}}]$:

$$\begin{aligned}
u_i(t) &= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [u_j(t - X)] \\
&= \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p_{ij}} [f(t - X - M_j)] \\
&\geq \max_{j \in \mathcal{V}(i)} \inf_{p_{ij} \in \mathcal{P}_{ij}} f(t - \mathbb{E}_{X \sim p_{ij}} [X] - M_j) \\
&\geq \max_{j \in \mathcal{V}(i)} f(t - m_{ij} - M_j) \\
&\geq f(t - \min_{j \in \mathcal{V}(i)} m_{ij} - M_j) \\
&\geq f(t - M_i) \\
&\geq \max_{j \in \mathcal{V}(i)} \mathbb{E}_{X \sim \delta_{m_{ij}}} [f(t - X - M_j)],
\end{aligned}$$

using the convexity and the monotonicity of $f(\cdot)$. The last inequality shows that $u_i(t) = f(t - M_i)$. This concludes the induction. Hence:

$$u_i(t) = f(t - M_i) \quad \forall t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\text{sup}}, T], \forall i \in \mathcal{V}.$$

Using Theorem 2 and plugging this last expression back into (7), we conclude that:

$$\pi_{f, \mathcal{P}}^*(i, t) \in \mathcal{T}^r(i) \quad \forall t \leq T, \forall i \neq d.$$

Moreover, for any arc $(i, j) \in \mathcal{A}$, we have proved that the infimum appearing in (7) is attained at $\delta_{m_{ij}}$ irrespective of the remaining budget t . This shows:

$$\begin{aligned}
\sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i, j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau} [f(T - X_\pi)] &\geq \sup_{\pi \in \Pi} \mathbb{E}_\delta [f(T - X_\pi)] \\
&\geq \sup_{\pi \in \Pi} \inf_{\forall (i, j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}} [f(T - X_\pi)],
\end{aligned}$$

where the notation δ refers to the fact the costs $(c_{ij})_{(i, j) \in \mathcal{A}}$ are independent and distributed according to $(\delta_{m_{ij}})_{(i, j) \in \mathcal{A}}$. Since (3) is a relaxation of (2), we get:

$$\sup_{\pi \in \Pi} \inf_{\forall (i, j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}} [f(T - X_\pi)] = \sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i, j) \in \mathcal{A}, p_{ij}^\tau \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}^\tau} [f(T - X_\pi)],$$

and an optimal strategy for both problems is to always follow the shortest-path tree \mathcal{T}^r . \square

B.8. Proof of Lemma 2

Proof of Lemma 2. Without loss of generality, we assume that $f^{(K+1)} > 0$. The proof is almost identical in the converse situation. We use Proposition 3 and consider a solution $(\pi_{f, \mathcal{P}}^*, (u_i(\cdot))_{i \in \mathcal{V}})$ to the dynamic programming equation (7). We first prove by induction on the level of the nodes i in \mathcal{G} that $u_i^{(K+1)} > 0$ and that, for j the immediate successor of i in \mathcal{G} , there exists $p_{ij} \in \mathcal{P}_{ij}$ such that:

$$u_i(t) = \mathbb{E}_{X \sim p_{ij}} [u_j(t - X)] \quad \forall t \leq T \text{ if } i \neq d. \quad (\text{EC.5})$$

Assume that the property holds for all nodes of level less than l and consider a node $i \in \mathcal{V}$ of level $l + 1$. Let j be the immediate successor of i in \mathcal{G} . As \mathcal{P}_{ij} is not empty, Lemma 3.1 from Shapiro (2001) shows that \mathcal{P}_{ij} contains a discrete distribution whose support is a subset of $\{\delta_0, \dots, \delta_{K+2}\}$ with $\delta_0 = \delta_{ij}^{\inf} < \delta_1 < \dots < \delta_{K+2} = \delta_{ij}^{\sup}$. For any $n \in \mathbb{N}$, we define the ambiguity set:

$$\mathcal{P}_{ij}^n = \{p \in \mathcal{P}_{ij} \mid \text{supp}(p) \subset \{\delta_0, \delta_0 + \frac{\delta_1 - \delta_0}{n}, \delta_0 + 2 \cdot \frac{\delta_1 - \delta_0}{n}, \dots, \delta_1, \delta_1 + \frac{\delta_2 - \delta_1}{n}, \dots, \delta_{K+1}\}\},$$

which can be interpreted as a discretization of \mathcal{P}_{ij} . Observe that, by design, \mathcal{P}_{ij}^n is not empty. Additionally, we define the sequence of functions $(f_i^n(\cdot))_{n \in \mathbb{N}}$ by:

$$f_i^n(t) = \inf_{p \in \mathcal{P}_{ij}^n} \mathbb{E}_{X \sim p}[u_j(t - X)] \quad \forall t \leq T. \quad (\text{EC.6})$$

Since $u_j^{(K+1)} > 0$, Prékopa (1990) shows that there exists $p_{ij}^n \in \mathcal{P}_{ij}^n$ such that:

$$f_i^n(t) = \mathbb{E}_{X \sim p_{ij}^n}[u_j(t - X)] \quad \forall t \leq T.$$

Because \mathcal{P}_{ij} is compact with respect to the weak topology and since $\mathcal{P}_{ij}^n \subset \mathcal{P}_{ij}$, we can take a subsequence of $(p_{ij}^n)_{n \in \mathbb{N}}$ such that $p_{ij}^n \rightarrow p_{ij} \in \mathcal{P}_{ij}$ as $n \rightarrow \infty$ for the weak topology. Without loss of generality, we continue to denote this sequence $(p_{ij}^n)_{n \in \mathbb{N}}$. Since $u_j(\cdot)$ is continuous, we derive that the sequence of functions $(f_i^n(\cdot))_{n \in \mathbb{N}}$ converges simply to a function $f_i(\cdot)$ which satisfies:

$$f_i(t) = \mathbb{E}_{X \sim p_{ij}}[u_j(t - X)] \quad \forall t \leq T. \quad (\text{EC.7})$$

We now move on to show that $f_i(t) = u_i(t)$ for all $t \leq T$. This will conclude the induction because we can take the $(K + 1)$ th derivative in (EC.7) since p_{ij} has compact support. Take $t \leq T$ and $\epsilon > 0$. The function $u_j(\cdot)$ is continuous on $[t - \delta^{\sup}, t - \delta^{\inf}]$ hence, by uniform continuity, there exists $\alpha > 0$ such that:

$$|u_j(t - \omega) - u_j(t - \omega')| \leq \epsilon$$

as soon as $|\omega - \omega'| \leq \alpha$ and $(\omega, \omega') \in [\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]^2$. Consider $n > \frac{\delta_{ij}^{\sup} - \delta_{ij}^{\inf}}{\alpha}$. Using conic duality, Corollary 3.1 of Shapiro (2001) shows that $u_i(t)$ is the optimal value of the infinite linear program:

$$\begin{aligned} & \sup_{(a_1, \dots, a_K, b) \in \mathbb{R}^{K+1}} \sum_{k=1}^K a_k \cdot m_{ij}^k + b \\ & \text{subject to} \quad \sum_{k=1}^K a_k \cdot \omega^k + b \leq u_j(t - \omega) \quad \forall \omega \in [\delta_{ij}^{\inf}, \delta_{ij}^{\sup}]. \end{aligned} \quad (\text{EC.8})$$

Using strong linear programming duality, we also have that $f_i^n(t)$ is the optimal value of the finite linear program:

$$\begin{aligned} & \sup_{(a_1, \dots, a_K, b) \in \mathbb{R}^{K+1}} \sum_{k=1}^K a_k \cdot m_{ij}^k + b \\ & \text{subject to} \quad \sum_{k=1}^K a_k \cdot \omega^k + b \leq u_j(t - \omega) \quad \forall \omega \in \{\delta_0, \delta_0 + \frac{\delta_1 - \delta_0}{n}, \delta_0 + 2 \cdot \frac{\delta_1 - \delta_0}{n}, \dots, \delta_{K+1}\}. \end{aligned} \quad (\text{EC.9})$$

Take $(a_1^n, \dots, a_K^n, b^n)$ an optimal basic feasible solution to (EC.9). By a standard linear programming argument:

$$\max\left(\max_{k=1, \dots, K} |a_k^n|, |b^n|\right) \leq U,$$

where $U = ((K+1) \cdot \max(1, u_j(t - \delta^{\inf}_{ij}), (\delta^{\sup}_{ij})^K))^{(K+1)}$ does not depend on n . Let us use the shorthand:

$$V = U \cdot (\delta^{\sup}_{ij} - \delta^{\inf}_{ij}) \cdot \sum_{k=1}^K k \cdot (\delta^{\sup}_{ij})^{(k-1)},$$

and define $b = b^n - \frac{V}{n} - \epsilon$. We show that (a_1^n, \dots, a_K^n, b) is feasible for (EC.8). For any $w \in [\delta^{\inf}_{ij}, \delta^{\sup}_{ij}]$, take $w' \in \{\delta_0, \delta_0 + \frac{\delta_1 - \delta_0}{n}, \delta_0 + 2 \cdot \frac{\delta_1 - \delta_0}{n}, \dots, \delta_{K+1}\}$ such that $|w - w'| \leq \frac{\delta^{\sup}_{ij} - \delta^{\inf}_{ij}}{n}$. We have:

$$\begin{aligned} \sum_{k=1}^K a_k^n \cdot \omega^k + b &= \sum_{k=1}^K a_k^n \cdot (\omega')^k + b^n + \sum_{k=1}^K a_k^n \cdot (\omega^k - (\omega')^k) - \frac{V}{n} - \epsilon \\ &\leq u_j(t - \omega') + \sum_{k=1}^K |a_k^n| \cdot |\omega^k - (\omega')^k| - \frac{V}{n} - \epsilon \\ &\leq u_j(t - \omega) + \sum_{k=1}^K U \cdot k \cdot (\delta^{\sup}_{ij})^{(k-1)} \cdot |\omega - \omega'| - \frac{V}{n} \\ &\leq u_j(t - \omega), \end{aligned}$$

where we use the fact that $(a_1^n, \dots, a_K^n, b^n)$ is feasible for (EC.9) in the first inequality, the uniform continuity of $u_j(\cdot)$ in the second and the definition of V in the last one. We derive:

$$f_i(t) - \frac{V}{n} - \epsilon \leq u_i(t) \leq f_i(t).$$

Taking $n \rightarrow \infty$ and $\epsilon \rightarrow 0$, we obtain $f_i(t) = u_i(t)$. This concludes the induction.

As a consequence of (EC.5), the infimum in (7) is always attain for p_{ij} , irrespective of the remaining budget t , so we can conclude that (2) and (3) are equivalent. \square

B.9. Proof of Lemma 3

This result is a direct consequence of the following observations:

- when the risk function is $f(t) = t$, following the shortest path with respect to $(\max_{p \in \mathcal{P}_{ij}} \mathbb{E}_{X \sim p}[X])_{(i,j) \in \mathcal{A}}$ is an optimal strategy for (3),
- when the risk function is $f(t) = \exp(t)$, following the shortest path with respect to $(\max_{p \in \mathcal{P}_{ij}} -\log(\mathbb{E}_{X \sim p}[\exp(-X)]))_{(i,j) \in \mathcal{A}}$ is an optimal strategy for (3),
- when the risk function is $f(t) = -\exp(-t)$, following the shortest path with respect to $(\max_{p \in \mathcal{P}_{ij}} \log(\mathbb{E}_{X \sim p}[\exp(X)]))_{(i,j) \in \mathcal{A}}$ is an optimal strategy for (3).

As a consequence, for any of these risk functions, (2) and (3) are equivalent. Define $g(\cdot)$ as any of these risk functions. Assuming that $\gamma \cdot g(t) + \beta \geq f(t) \geq a \cdot g(t) + b, \forall t \leq T$, we get:

$$\sup_{\pi \in \Pi} \inf_{\forall (i,j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}}[f(T - X_{\pi})] \leq \sup_{\pi \in \Pi} \inf_{\forall (i,j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{P}}[\gamma \cdot g(T - X_{\pi}) + \beta]$$

$$\begin{aligned}
&\leq \beta + \gamma \cdot \sup_{\pi \in \Pi} \inf_{\forall (i,j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}}[g(T - X_{\pi})] \\
&\leq \beta + \gamma \cdot \sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i,j) \in \mathcal{A}, p_{ij}^{\tau} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^{\tau}}[g(T - X_{\pi})] \\
&\leq \beta - \frac{\gamma}{a} \cdot b + \frac{\gamma}{a} \cdot \sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i,j) \in \mathcal{A}, p_{ij}^{\tau} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^{\tau}}[a \cdot g(T - X_{\pi}) + b] \\
&\leq \beta - \frac{\gamma}{a} \cdot b + \frac{\gamma}{a} \cdot \sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i,j) \in \mathcal{A}, p_{ij}^{\tau} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^{\tau}}[f(T - X_{\pi})].
\end{aligned}$$

This last inequality along with:

$$\sup_{\pi \in \Pi} \inf_{\forall (i,j) \in \mathcal{A}, p_{ij} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}}[f(T - X_{\pi})] \geq \sup_{\pi \in \Pi} \inf_{\forall \tau, \forall (i,j) \in \mathcal{A}, p_{ij}^{\tau} \in \mathcal{P}_{ij}} \mathbb{E}_{\mathbf{p}^{\tau}}[f(T - X_{\pi})]$$

yields the claim with some basic algebra.

B.10. Proof of Lemma 4

Proof of Lemma 4. For any $k \in \mathbb{N}$, we define $(\pi^k, (u_i^k(\cdot))_{i \in \mathcal{V}})$ as a solution to the dynamic program (7) when the ambiguity sets are taken as $(\mathcal{P}_{ij}^k)_{(i,j) \in \mathcal{A}}$. Similarly, we define $(\pi^{\infty}, (u_i^{\infty}(\cdot))_{i \in \mathcal{V}})$ as a solution to the dynamic program (7) when the ambiguity sets are taken as $(\cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k)_{(i,j) \in \mathcal{A}}$. Along the same lines as what is done in the proof of Proposition 4, we can show that the functions $(u_i^k(\cdot))_{k \in \mathbb{N}}$ and $u_i^{\infty}(\cdot)$ are continuous for any $i \in \mathcal{V}$. Because the ambiguity sets are nested, observe that the sequence $(u_i^k(t))_{k \in \mathbb{N}}$ is non-decreasing for any $t \leq T$, hence it converges to a limit $f_i(t) \leq u_i^{\infty}(t)$. Moreover, $f_d(t) = f(t)$ for all $t \leq T$. Take $i \neq d$ and $t \leq T$. We have, for any $k \in \mathbb{N}$ and $m \leq k$:

$$\begin{aligned}
f_i(t) &\geq u_i^k(t) \\
&\geq \max_{j \in \mathcal{V}(i)} \inf_{p \in \mathcal{P}_{ij}^k} \int_0^{\infty} p(\omega) \cdot u_j^k(t - \omega) d\omega \\
&\geq \max_{j \in \mathcal{V}(i)} \int_0^{\infty} p_{ij}^k(\omega) \cdot u_j^k(t - \omega) d\omega \\
&\geq \max_{j \in \mathcal{V}(i)} \int_0^{\infty} p_{ij}^k(\omega) \cdot u_j^m(t - \omega) d\omega,
\end{aligned}$$

where $p_{ij}^k \in \mathcal{P}_{ij}^k$ achieves the minimum for any $j \in \mathcal{V}(i)$, which can be shown to exist since \mathcal{P}_{ij}^k is compact and $u_j^k(\cdot)$ is continuous. Because \mathcal{P}_{ij}^k is compact for the weak topology, we can take a subsequence of $(p_{ij}^k)_{k \in \mathbb{N}}$ that converges to a distribution p_{ij}^{∞} in $\cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k$. Without loss of generality we continue to refer to this sequence as $(p_{ij}^k)_{k \in \mathbb{N}}$. Taking the limit $k \rightarrow \infty$ in the last inequality derived yields:

$$f_i(t) \geq \max_{j \in \mathcal{V}(i)} \int_0^{\infty} p_{ij}^{\infty}(\omega) \cdot u_j^m(t - \omega) d\omega.$$

Observing that $u_j^m(t - \omega) \geq u_j^1(t - \omega)$, we can use the monotone convergence theorem for $m \rightarrow \infty$ and conclude that:

$$\begin{aligned}
f_i(t) &\geq \max_{j \in \mathcal{V}(i)} \int_0^{\infty} p_{ij}^{\infty}(\omega) \cdot f_j(t - \omega) d\omega \\
&\geq \max_{j \in \mathcal{V}(i)} \inf_{p \in \cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k} \int_0^{\infty} p(\omega) \cdot f_j(t - \omega) d\omega.
\end{aligned}$$

We use Theorem 2 for the ambiguity sets $(\cap_{k \in \mathbb{N}} \mathcal{P}_{ij}^k)_{(i,j) \in \mathcal{A}}$ and denote by T_f^r (resp. \mathcal{T}^r) the time budget (resp. the tree) put forth in the statement of the theorem. Using the last sequence of inequalities derived, we can prove, by induction on the levels of the nodes in \mathcal{T}^r that:

$$f_i(t) \geq u_i^\infty(t) \quad \forall t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\text{sup}}, T_f^r], \forall i \in \mathcal{V},$$

and then by induction on $m \in \mathbb{N}$ that:

$$f_i(t) \geq u_i^\infty(t) \quad \forall t \in [T_f^r - (|\mathcal{V}| - \text{level}(i, \mathcal{T}^r) + 1) \cdot \delta^{\text{sup}}, T_f^r + m \cdot \delta^{\text{inf}}], \forall i \in \mathcal{V}.$$

We finally obtain $f_s(T) \geq u_s^\infty(T)$ which concludes the proof.

□

B.11. Proof of Lemma 7

Proof of Lemma 7. First observe that, along the same lines as in general case, the constraint

$$z + (x - y) \cdot l \cdot \Delta t \leq u_j^{\Delta t}((k - l) \cdot \Delta t)$$

does not limit the feasible region if $(l \cdot \Delta t, u_j^{\Delta t}((k - l) \cdot \Delta t))$ is not an extreme point of the upper convex hull of $\{(l \cdot \Delta t, u_j^{\Delta t}(l \cdot \Delta t)), l = k - \left\lceil \frac{\delta_{ij}^{\text{sup}}}{\Delta t} \right\rceil, \dots, k - \left\lfloor \frac{\delta_{ij}^{\text{inf}}}{\Delta t} \right\rfloor\} \cup \{(\delta_{ij}^{\text{sup}}, u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\text{sup}})), (\delta_{ij}^{\text{inf}}, u_j^{\Delta t}(k \cdot \Delta t - \delta_{ij}^{\text{inf}}))\}$. Hence, we can discard the constraints that do not satisfy this property from (14). We denote by S the sorted projection of the set of extreme points onto the first coordinate. Observe that the feasible region is pointed as the polyhedron described by the inequality constraints does not contain any line, therefore there exists a basic optimal feasible solution for which at least three inequality constraints are binding. By definition of S , only two of the constraints

$$z + (x - y) \cdot \omega \leq u_j^{\Delta t}(\omega) \quad \omega \in S$$

can be binding which further implies that at least one of the constraints $x \geq 0$ and $y \geq 0$ must be binding. There are three types of feasible basis depending on whether these last two constraints are binding or if only one of them is. We show that, for each type, we can identify an optimal basis among the basis of the same type by binary search on the first coordinate of the extreme points. This will conclude the proof as it takes constant time to compare the objective function achieved by each of the three potentially optimal basis. Since, by definition of S , $u_j^{\Delta t}(\cdot)$ is convex on S , we can partition S into S_1 and S_2 such that $u_j^{\Delta t}(\cdot)$ is non-increasing on S_1 and non-decreasing on S_2 with $\max(S_1) = \min(S_2)$.

If $x \geq 0$ and $y \geq 0$ are binding then z is the only non-zero variable and the objective is to maximize z . Hence, the optimal basis of this type is given by $x = 0$, $y = 0$ and $z = \min_{\omega \in S} u_j^{\Delta t}(\omega)$ which can be computed by binary search since $u_j^{\Delta t}(\cdot)$ is convex on S .

If only $x \geq 0$ is binding, then the line $\omega \rightarrow z - y \cdot \omega$ must be joining two consecutive points in S_1 . Since the objective function is precisely the value taken by the line $\omega \rightarrow z - y \cdot \omega$ at β^{ij} , the optimal straight line joins two consecutive points in S_1 , ω_1 and ω_2 , that satisfy $\omega_1 \leq \beta^{ij} \leq \omega_2$ assuming $\max(S_1) \geq \beta^{ij}$. If $\max(S_1) < \beta^{ij}$, the feasible basis of this type are dominated by the optimal basis of the first type. Computing ω_1 and ω_2 or showing that they do not exist can be done with a single binary search on S .

The discussion is analogous if only $y \geq 0$ is binding instead. The line $\omega \rightarrow z + x \cdot \omega$ must be joining two consecutive points in S_2 . Since the objective function is precisely the value taken by this line at α^{ij} , the optimal straight line joins two consecutive points in S_2 , ω_1 and ω_2 , that satisfy $\omega_1 \leq \alpha^{ij} \leq \omega_2$ assuming $\alpha^{ij} \geq \min(S_2)$. If $\min(S_2) > \alpha^{ij}$, the feasible basis of this type are dominated by the optimal basis of the first type. Computing ω_1 and ω_2 or showing that they do not exist can be done with a single binary search on S .

□