

Fare Evasion in Transit Networks

 José Correa,^a Tobias Harks,^b Vincent J. C. Kreuzen,^c Jannik Matuschke^d
^aDepartamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile; ^bInstitut für Mathematik, Universität Augsburg, 86159 Augsburg, Germany; ^cSchool of Business and Economics, Maastricht University, 6211 LK, Maastricht, The Netherlands;

^dTUM School of Management, Technische Universität München, 80333 München, Germany

Contact: correa@uchile.cl (JC); tobias.harks@math.uni-augsburg.de (TH); v.kreuzen@maastrichtuniversity.nl (VJCK); jannik.matuschke@tum.de (JM)

Received: September 16, 2014

Revised: May 4, 2016; July 15, 2016

Accepted: August 5, 2016

Subject Classifications: networks/graphs; mathematics: combinatorics; industries: transportation/shipping

Area of Review: Transportation

<https://doi.org/10.1287/opre.2016.1560>
Copyright: © 2017 INFORMS

Abstract. Public transit systems in major urban areas usually operate under deficits and therefore require significant subsidies. An important cause of this deficit, particularly in the developing world, is the high fare evasion rate mainly due to an ineffective control policy or the lack of it. In this paper we study new models for optimizing fare inspection strategies in transit networks based on bilevel programming. In the first level, the leader (the network operator) determines probabilities for inspecting passengers at different locations, while in the second level, the followers (the fare-evading passengers) respond by optimizing their routes given the inspection probabilities and travel times. To model the followers' behavior we study both a nonadaptive variant, in which passengers select a path a priori and continue along it throughout their journey, and an adaptive variant, in which they gain information along the way and use it to update their route. For these problems, which are interesting in their own right, we design exact and approximation algorithms, and we prove a tight bound of $3/4$ on the ratio of the optimal cost between adaptive and nonadaptive strategies. For the leader's optimization problem, we study a fixed-fare and a flexible-fare variant, where ticket prices may or may not be set at the operator's will. For the latter variant, we design an LP-based approximation algorithm. Finally, employing a local search procedure that shifts inspection probabilities within an initially determined support set, we perform an extensive computational study for all variants of the problem on instances of the Dutch railway and the Amsterdam subway network. This study reveals that our solutions are within 5% of theoretical upper bounds drawn from the LP relaxation. We also derive exact nonlinear programming formulations for all variants of the leader's problem and use them to obtain exact solutions for small instance sizes.

Funding: This work was supported by the Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 and by the German Academic Exchange Service (DAAD).

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/opre.2016.1560>.

Keywords: transit networks • bilevel optimization

1. Introduction

Fare evasion in transit systems causes significant operational deficits that have to be compensated by large subsidies. The issue is particularly relevant in the developing world where evasion rates are often very high. In Santiago de Chile, for instance, evasion rates of nearly 20% for the public transportation system, Transantiago, lead to an estimated annual cost of 450 million dollars (Torres-Montoya 2014). This accounts for a large fraction of the total deficit of the system estimated at 700 million dollars each year (Gómez-Lobo 2012, Chilean Subsecretary of Transportation 2014). Although the situation is less dramatic in the developed world, fare evasion is still a major source of inefficiencies. For example, recent studies revealed an annual loss of 70 million pounds for London's transit system (Transport for London 2010).

As the installation of physical ticket barriers is not always possible, cost-efficient, or desirable, many

transit systems, such as the Dutch and German railway and subway networks and the London and Santiago bus systems, rely on the honesty of customers and proper controlling by ticket inspectors on board. Travelers who are caught without a valid proof of payment have to pay a fine, which is significantly larger than the ticket price.

In this paper we study the optimization of fare inspection strategies in transit systems taking into account realistic models for the passenger's reaction. Our models are based on a *bilevel optimization problem* (or Stackelberg game) on the network. In the first level, the leader, who strives to maximize the revenue from ticket sales and collected fines, determines for each edge the probability of controlling passing passengers, representing the frequency of inspections on that edge. In light of the limited number of inspectors available, we assume a global budget constraint on the sum of these probabilities. Given the inspection probabilities

on edges, we assume passengers act strategically and decide on whether or not to buy a ticket and along which path to travel based solely on their perceived cost. If a passenger chooses not to buy a ticket, her cost is thus the sum of the path length (expressed in monetary units) and the expected fine to be payed (which depends on the inspection probabilities along the path); on the other hand, if she buys a ticket, her cost is defined as the cost of the cheapest path considering both ticket price and distance measured in monetary units.

We perform a complete study differentiating between two possibilities for the passengers' behavioral assumption as well as two settings for the leader's decision problem. From the followers' side, we consider both the case in which they are *adaptive* and *nonadaptive*. In the latter, each passenger chooses a path at the beginning of her journey and continues along it, independent of whether or not she encounters an inspector. In the adaptive version, passengers adapt their behavior and consequently, when caught without a ticket, they continue on their shortest path only considering distance (since the fine typically includes a ticket to finish the trip). For the leader, we also study two different settings. In the *fixed-fare* setting, we assume that ticket prices are fixed a priori, e.g., by governmental regulations, and the leader only sets the inspection probabilities subject to a budget constraint. In the *flexible-fare* setting, we assume that the leader can additionally determine ticket prices.

Our Contribution and Structure of the Paper. The main goal of this work is to provide a comprehensive study of fare evasion and fare inspection problems in transit networks. As mentioned above, we consider four versions of the problem by varying on whether the followers are adaptive or not, and whether the ticket prices are fixed or flexible. From a methodological viewpoint we tackle the problem from different angles, designing polynomial time algorithms and approximation schemes for the followers' problems, studying approximation algorithms and LP relaxations for the leader's problem, and using a local search heuristic to obtain high quality solutions for real-world instances.

After establishing the precise model in Section 2, we study the two variants of the followers' minimization problem in Section 3. These are natural extensions of the classic shortest path problem and exhibit interesting properties in their own right. For the nonadaptive variant, we design a fully polynomial time approximation scheme for general network topologies, exploiting similarities between this problem and the restricted shortest path problem. We also establish a close connection between the nonadaptive followers' problem and the parametric shortest path problem, showing that the optimal solution to the former can be found by enumerating all pieces of the optimal cost curve of the

parametric problem. Although for general graphs the number of pieces can be superpolynomial in the size of the network, we show that for series-parallel graphs the number is bounded by the number of edges, supporting a conjecture by Nikolova (2009) claiming the same bound for planar graphs. For adaptive followers, we obtain an exact polynomial time algorithm for general graphs using an optimal substructure property of the problem, and we also present an exact LP formulation that reveals an interesting relation between the adaptive follower's problem and generalized flows. Surprisingly, we further show that for arbitrary probability distributions, the optimal solution to the nonadaptive followers' minimization problem is at most a factor $4/3$ away from the optimal solution to the adaptive variant, and we also show that this bound is tight.

In Section 4 we turn our attention to the leader's maximization problem and prove that all four variants of this problem are strongly NP-hard. We then present an LP relaxation that yields a valid upper bound on the achievable profit for all four variants, and we also obtain a $(1 - 1/e)$ -approximation algorithm for the variant involving flexible ticket prices and nonadaptive followers. Combining this with the worst-case gap from the preceding section yields a $\frac{3}{4}(1 - 1/e)$ -approximation for the variant with adaptive followers. We further derive exact nonlinear formulations for all four different versions of the leader's problem, using the generalized flow LP from the preceding section. We then present a local search procedure that shifts inspection probabilities within an initially determined support set of edges. As candidate support sets, we use solutions from an LP relaxation, a minimum multicut, and a related mixed integer program proposed by Borndörfer et al. (2013). Finally, we discuss that all results in this section also work in the presence of elastic demands, i.e., when followers have an outside option to avoid using the transit network if costs for doing so are too high.

In Section 5 we demonstrate the applicability of our local search approach for all four problem variants by conducting a computational study comprising a total of 5,600 instances based on the networks of the Dutch railway system, the Amsterdam metro system, and randomly generated graphs. Our study reveals that the objective values of the computed solutions are, on average, within 5% of the calculated upper bounds. For several small instances, we even obtain exact solutions (or improved upper bounds) by solving the exact (but nonconvex) formulations from Section 4 using the global optimization solver BARON (Tawarmalani and Sahinidis 2005). It turns out that for those instances for which BARON obtained provably optimal solutions, our solutions (computed with the local search procedure) are, in fact, within 2% of optimality.

Related Work. Bilevel network pricing problems have been considered extensively in the literature. In the work of Brotcorne et al. (2001, 2008) and Labbé et al. (1998) the leader prices edges of some subgraph and the followers choose shortest paths with respect to edge costs defined as the sum of travel costs and prices. This problem was proved to be NP-hard by Labbé et al. (1998) and later several hardness of approximation and approximability results have been obtained (Bouhtou et al. 2007, Briest et al. 2012, Roch et al. 2005).

Borndörfer et al. (2013, 2012a) investigated a problem similar to our variant of nonadaptive followers and fixed ticket prices. In contrast to our model, they assume the cost of the followers to be given by the sum of travel costs and probabilities, reducing the followers' minimization problem to a standard shortest path problem. Based on this simplification, the authors derive a compact mixed integer program (MIP) for the leader's optimization problem and solve it on several test instances derived from the German motorway network. In addition, they also study the case where the leader wants to minimize the number of fare evaders. In contrast, we model the followers' response more precisely by considering the exact expected fine along a path, leading to a nonlinear aggregation of probabilities. In our computational study, we systematically compare the difference between both modeling approaches. We show that the quality of plain MIP solutions is, on average, over all instances 5% lower than our solutions when taking the exact followers' response into account. For instances derived from the Amsterdam metro network, our solution improves even by 7.4% on average, with an improvement of more than 20% on some instances.

There is a body of literature on so-called network security games, where a defender can check a limited number of edges in a graph and the attacker computes from a given set of source and target nodes a path connecting a source and a target (Jain et al. 2013, Tambe 2012, Washburn and Wood 1995). The payoff for the attacker (and defender) has a binary characteristic since it only depends on whether or not the attacker has been caught. Lin et al. (2013) discuss a model in which a single defender moves through the network over time to detect attackers. The authors propose a linear program—which solves the problem to optimality for small instances—and heuristics—which solve the problem for larger instances within 1% of optimality.

Yin et al. (2012) study a fare inspection model similar to ours, although, in contrast to our work, passengers are assumed to follow a fixed route through the network. On the other hand, Borndörfer et al. (2012a) consider the problem of computing tours of inspectors for a given distribution of traffic flows. They derive an integer programming formulation and conduct a computational study on the German motorway network.

2. The Model

We are given a directed graph $G = (V, E)$ with costs $c: E \rightarrow \mathbb{Z}_+$ modeling the transit times on the edges (as monetary cost incurred to the passengers traveling along them). To prevent fare evasion, the network operator (the *leader*) sets inspection probabilities $p_e \in [0, 1]$ on the edges subject to the budget constraint $\sum_{e \in E} p_e \leq B$, where $B \geq 0$ corresponds to a limited number of ticket inspectors. We first describe the passengers' reaction to the chosen inspection strategy and then discuss the resulting revenue for the operator.

2.1. The Followers

Passengers (the *followers*) are modeled by a set of commodities K . For every commodity $i \in K$, a demand $d_i \geq 0$ is given that specifies the number of passengers with origin $s_i \in V$ and destination $t_i \in V$. Passengers of commodity i can either buy a ticket at price $T_i \geq 0$ or choose a path without paying the ticket. If a passenger is caught without a valid ticket, he has to pay a fine $F \geq 0$ satisfying $F \geq T_i$. As in most public transport systems (e.g., the Dutch or German subway and railway networks), the fine includes the ticket price and enables the passenger to continue his trip. Passengers are assumed to be rational, deciding purely based on their personal costs, which is expressed as the sum of travel times along the chosen path and the monetary cost. In the following, we differentiate between a *nonadaptive* and an *adaptive* variant of the followers' response. In both settings, passengers of commodity i can either choose to pay the ticket and follow the shortest path w.r.t. c , or decide to evade the fare, choosing a path $P \in \mathcal{P}_i$ that minimizes a variant-dependent cost function (where \mathcal{P}_i denotes the set of s_i - t_i -paths).

Nonadaptive followers. In the nonadaptive variant, passengers are assumed to choose a route before the start of their trip and continue along it, independent of whether or not they encounter an inspector. For a path P , we denote its total travel time by $c(P) := \sum_{e \in P} c_e$, and we denote the probability that no inspector is encountered along P by $\pi(P) := \prod_{e \in P} (1 - p_e)$. Given probabilities p , the expected cost of path P for passenger i in the nonadaptive setting is

$$f_{N,p,i}(P) := c(P) + (1 - \pi(P)) \cdot F.$$

This variant is plausible under the assumption that passengers have determined their route beforehand (via checking the route and timetables). We denote the corresponding optimization problem by FMP_N .

Adaptive followers. In the adaptive variant, once a passenger is caught, he will continue his trip along

the shortest path (w.r.t. c). Letting $SP_c(v, w)$ denote the shortest path distance from v to w w.r.t. c , the follower thus tries to find a path $P = (e_1, \dots, e_l) \in \mathcal{P}_i$ with $e_j = (v_j, v_{j+1})$ minimizing the expected cost

$$f_{A,p,i}(P) := \sum_{j=1}^l \prod_{k=1}^{j-1} (1 - p_{e_k}) \cdot (c_{e_j} + p_{e_j}(F + SP_c(v_{j+1}, t_i))).$$

Note that in this formula, the j th summand corresponds to the event of v_j being reached without inspection, in which case e_j is traversed next. This variant assumes that passengers know the shortest paths from all stations to their destination in a network, e.g., by calculating it using a portable computing device. We denote the corresponding optimization problem by FMP_A .

Throughout the paper, we will make use of the following lower bounds on $f_{A,p,i}$ and $f_{N,p,i}$. Intuitively, these bounds follow from the fact that a passenger always has to traverse a path that is at least as long as the shortest path from s_i to t_i .

Lemma 1. *Let $i \in K$ and $P \in \mathcal{P}_i$. Then $f_{A,p,i}(P) \geq \pi(P)c(P) + (1 - \pi(P)) \cdot (SP_c(s_i, t_i) + F)$. Furthermore, if $c(P) = SP_c(s_i, t_i)$ then the above inequality holds with equality.*

Proof. Define $\pi_j := \prod_{k=1}^j (1 - p_{e_k})$. Note that $\pi_{j-1}p_{e_j} = \pi_{j-1}(p_{e_j} - 1 + 1) = \pi_{j-1} - \pi_j$. We obtain

$$\begin{aligned} f_{A,p,i}(P) &= \sum_{j=1}^l \pi_{j-1} \cdot (c_{e_j} + p_{e_j}(F + SP_c(v_{j+1}, t_i))) \\ &= \sum_{j=1}^l \pi_{j-1} \cdot \left(c_{e_j} - p_{e_j} \sum_{k=1}^j c_{e_k} \right. \\ &\quad \left. + p_{e_j} \left(F + \underbrace{SP_c(v_{j+1}, t_i) + \sum_{k=1}^j c_{e_k}}_{\geq SP_c(s_i, t_i)} \right) \right) \\ &\geq \sum_{j=1}^l \pi_{j-1} \left(c_{e_j} - p_{e_j} \sum_{k=1}^j c_{e_k} \right) + \sum_{j=1}^l \pi_{j-1} p_{e_j} (F + SP_c(s_i, t_i)) \\ &= \sum_{j=1}^l \left(\pi_{j-1} - \sum_{k=j}^l \pi_{k-1} p_{e_k} \right) c_{e_j} + \sum_{j=1}^l (\pi_{j-1} - \pi_j) \\ &\quad \cdot (F + SP_c(s_i, t_i)) \\ &= \sum_{j=1}^l \left(\pi_{j-1} - \sum_{k=j}^l (\pi_{k-1} - \pi_k) \right) c_{e_j} \\ &\quad + (1 - \pi_l)(F + SP_c(s_i, t_i)) \\ &= \pi_l \cdot c(P) + (1 - \pi_l)(F + SP_c(s_i, t_i)). \end{aligned}$$

Note that the only inequality used above is $SP_c(v_{j+1}, t_i) + \sum_{k=1}^j c_{e_k} \geq SP_c(s_i, t_i)$, which holds with equality for every j if P is a shortest path with respect to c . \square

Corollary 1. *Let $X \in \{A, N\}$, $i \in K$, and $P \in \mathcal{P}_i$. Then $f_{X,p,i}(P) \geq SP_c(s_i, t_i) + (1 - \pi(P)) \cdot F$.*

2.2. The Leader

The leader’s problem can be defined as a bilevel problem, where the leader sets probabilities on the edges to which the followers respond by solving their individual optimization problems. While we will always assume the fine F to be fixed in the problem input (fines are commonly determined by the legislation (Gijsbers 2013)), we will consider both the scenario where ticket prices are *fixed* a priori and the scenario where they are *flexible* and can be determined by the leader as part of his optimization problem. Combining this with the two different models for the followers’ reaction, we obtain four different versions of the *leader’s maximization problem*, which we will denote by LMP_X^L , with $L \in \{\text{fix}, \text{flex}\}$ and $X \in \{N, A\}$ specifying the ticket pricing variant and the behaviour of the followers, respectively. Accordingly, we denote by $\Gamma_{X,i}^L(p)$ the leader’s revenue per passenger received from commodity i when choosing inspection probabilities p . Thus, the leader wants to maximize $\sum_{i \in K} d_i \Gamma_{X,i}^L(p)$.

Fixed Fares. In the fixed-fare setting, the revenue received from passenger i is either the ticket price or the expected revenue from collecting fines, i.e.,

$$\Gamma_{X,i}^{\text{fix}}(p) := \max \left\{ F \cdot (1 - \pi(P^*)) : P^* \in \arg \min_{P \in \mathcal{P}_i \cup \{P_i^T\}} f_{X,p,i}(P) \right\},$$

where P_i^T denotes a special path representing the option of paying the ticket with $f_{X,p,i}(P_i^T) := SP_c(s_i, t_i) + T_i$ and $\pi(P_i^T) := 1 - T_i/F$.

Flexible Fares. When the leader is allowed to determine the ticket prices, the optimal choice given the probabilities p is to set $T_i = \min\{f_{X,p,i}(P) : P \in \mathcal{P}_i\} - SP_c(s_i, t_i)$ for every $i \in K$. To see this, observe that this is the maximum ticket price that follower i is willing to pay, because for this choice of T_i the cost of paying the ticket and traveling along the shortest path is equal to the expected cost of the best fare evasion option. Thus, decreasing the ticket price would not change the follower’s decision to pay the ticket and only lower the leader’s revenue. Increasing the ticket price would result in the follower not paying the ticket but choosing a fare evasion route $P^* \in \mathcal{P}_i$ with $f_{X,p,i}(P^*) = \min_{P \in \mathcal{P}_i} f_{X,p,i}(P)$. Note that $f_{X,p,i}(P^*) \geq (1 - \pi(P^*))F + SP_c(s_i, t_i)$ by Corollary 1. Hence, the revenue received from the fare evading passenger is $(1 - \pi(P^*))F \leq f_{X,p,i}(P^*) - SP_c(s_i, t_i) = T_i$ in this case. We conclude that the optimal revenue received from passenger i in the flexible fare setting is indeed achieved by setting the tickets prices as described above; thus, we define

$$\Gamma_{X,i}^{\text{flex}}(p) := \min\{f_{X,p,i}(P) : P \in \mathcal{P}_i\} - SP_c(s_i, t_i).$$

3. The Followers' Minimization Problem

In this section we design efficient algorithms of both variants of the followers' minimization problem. Throughout this section we assume we are given the graph $G = (V, E)$, the start $s \in V$ and destination $t \in V$ of a given follower, costs c , probabilities p , and the fine F .

3.1. Nonadaptive Followers' Minimization Problem

We first turn our attention to the nonadaptive version of the followers' minimization problem FMP_N and derive a fully polynomial-time approximation scheme (FPTAS), a polynomial time exact algorithm for series-parallel graphs, and a nonpolynomial exact algorithm for general graphs.

A Fully Polynomial Time Approximation Scheme. The nonadaptive version of the followers' minimization problem is related to the *restricted shortest path problem* (RSP): If P^* is an s - t -path minimizing f_N , then P^* also maximizes $\pi(P^*)$ among all paths with $c(P) \leq C^* := c(P^*)$. By discretizing the set of possible values for C^* , this relation can be used to derive an FPTAS for FMP_N .

Theorem 1. *There is an algorithm for FMP_N that computes a $(1 + \varepsilon)$ -approximate solution in time polynomial in $1/\varepsilon$, $\log c_{\max}$, $|V|$, and $|E|$, where $c_{\max} := \max_{e \in E} c_e$.*

Proof. Define $k := \lceil \log_{1+\varepsilon} |V| c_{\max} \rceil$, and $C_i := (1 + \varepsilon)^i$ for $i \in [k]$. For each $i \in [k]$, we use the restricted shortest path algorithm by Lorenz and Raz (2001) to compute a path P_i with $c(P_i) \leq (1 + \varepsilon)C_i$ and $\pi(P_i) \geq \pi(P)$ for all P with $c(P) \leq C_i$. Choose a path P' from P_0, \dots, P_k minimizing $f_N(P')$ among all s - t -paths of the graph. Let P^* be an s - t -path minimizing f_N and define $i^* := \min\{i \in [k] : c(P^*) \leq C_i\}$. Then $c(P_{i^*}) \leq (1 + \varepsilon)^2 c(P^*)$ and $\pi(P_{i^*}) \geq \pi(P^*)$ and thus $f_N(P') \leq f_N(P_{i^*}) \leq (1 + \varepsilon)^2 f_N(P^*)$. \square

Exact Algorithms Using Parametric Shortest Paths. In the *parametric shortest path problem*, the cost of each arc e is specified by an affine linear function $c_e(\lambda) = a_e \lambda + b_e$ of a parameter λ . The goal is to find a set of paths that contains an optimal path for each choice of λ . Given an instance of FMP_N , we will construct an instance of the parametric shortest path problem with the property that there exists a value λ^* of the parameter such that any path that is optimal for λ^* is also an optimal solution to the original FMP_N instance.

First note that without loss of generality, we can assume $p_e < 1$ for all $e \in E$: If an optimal path P^* contains an arc e with $p_e = 1$, then $\pi(P^*) = 0$ and, thus, $f_N(P^*) = c(P^*) + F$. Optimality of P^* then implies $c(P^*) \leq f_N(P) - F = c(P) - \pi(P)F$ for all $P \in \mathcal{P}$, and hence P^* must be a shortest path with respect to c . We can, therefore, solve the problem separately for the graph in which all arcs with $p_e = 1$ are removed and later compare the solution to a shortest path with respect to c .

Now for every $e \in E$, define $q_e := -\log(1 - p_e)$ and $g_\lambda(P) := c(P) + \lambda q(P)$. We will make use of the following observation, which is a variation of a lemma proven by Nikolova et al. (2006).

Lemma 2. *There is a $\lambda^* \geq 0$ such that any $P \in \mathcal{P}$ that minimizes $g_{\lambda^*}(P)$ also minimizes $f_N(P)$.*

Proof. Consider the polyhedron $\mathcal{Y} := \text{conv}\{(c(P), q(P)) : P \in \mathcal{P}\} + \mathbb{R}_+^2$ and the function $h: \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by $h(c, q) := c + (1 - 2^{-q})F$. Observe that h is concave and component-wise nondecreasing. Therefore, h attains its minimum over \mathcal{Y} at a vertex (c^*, q^*) of \mathcal{Y} . Let $\lambda^* \geq 0$ be such that (c^*, q^*) is the unique optimal solution to $\min_{(c, q) \in \mathcal{Y}} c + \lambda^* q$. Let P^* be a path minimizing $g_{\lambda^*}(P)$. This implies $c(P^*) = c^*$ and $q(P^*) = q^*$. Therefore, P^* also minimizes $h(c(P), q(P)) = f_N(P)$. \square

Although we do not know the correct value of λ^* a priori, Lemma 2 implies that we can find an optimal solution to FMP_N by constructing the complete optimal cost curve of the parametric shortest path problem. This curve is a piecewise linear function that can be constructed by parametric search solving at most $O(k)$ standard shortest path problems (Carstensen 1983), where k is number of breakpoints of the curve. For completeness, we list a simple algorithm for constructing the curve and computing the optimal path as Algorithm 1. For general graphs, it is known that the number of pieces is bounded by $|V|^{\log|V|+1}$ and that this bound is tight; see the original result by Carstensen (1983) and the simplified proof by Nikolova (2009). We can, therefore, solve FMP_N in quasi-polynomial time.

Algorithm 1 (Algorithm for FMP_N using the parametric shortest path problem)

```

Let  $P_1$  be a path minimizing  $q(P_1)$ .
Let  $P_2$  be a path minimizing  $c(P_2)$ .
if  $q(P_1) = q(P_2)$  then
    | return  $P_2$ 
else
    |  $\mathcal{P} \leftarrow \{P_1, P_2\} \cup \text{FINDBREAKPOINTS}(P_1, P_2)$ 
    | Choose  $P' \in \arg \min_{P \in \mathcal{P}} f_N(P)$ .
    | return  $P'$ 
procedure  $\text{FINDBREAKPOINTS}(P_1, P_2)$ 
    |  $\lambda \leftarrow \frac{c(P_1) - c(P_2)}{q(P_2) - q(P_1)}$ 
    | Let  $P$  be a path minimizing  $g_\lambda(P)$ .
    | if  $g_\lambda(P) = g_\lambda(P_1)$  then
    | | return  $\emptyset$ 
    | else
    | | return  $\text{FINDBREAKPOINTS}(P_1, P) \cup \{P\}$ 
    | |  $\cup \text{FINDBREAKPOINTS}(P, P_2)$ 
    
```

Theorem 2. *There is an algorithm that solves FMP_N in time $O(|V|^{\log|V|})$.*

Nikolova (2009) conjectured that for the case of planar graphs the number of breakpoints is polynomial in the number of edges. We show that this conjecture is true for series-parallel graphs. This class of graphs is of particular interest as most hardness reductions for shortest paths problems with two types of cost, such as the restricted shortest path problem, are based on a special subclass of series-parallel graphs; see, e.g., the reduction by Klinz and Woeginger (2004) for minimum cost flows over time in series-parallel graphs.

Theorem 3. *For series-parallel graphs, the optimal cost curve of the parametric shortest path problem has at most $|E|$ pieces.*

Proof. We prove the statement via induction over the construction of a series-parallel graph as the series or parallel composition of two smaller series-parallel graphs. The base case is a graph consisting of a single arc, in which case the claim is trivially true.

Now consider the case that G is the series composition of two series-parallel graphs $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$ joined at a vertex v . Then for every value of λ , the path minimizing $c_\lambda(P)$ is the concatenation of an s - v -path P_1 in G_1 and a v - t -path P_2 in G_2 , each minimizing c_λ . Therefore, each breakpoint of the optimal cost curve for G is also a breakpoint of the curve for G_1 or G_2 . By induction hypothesis, there are at most $|E_1| - 1$ breakpoints for G_1 and at most $|E_2| - 1$ breakpoints for G_2 . This implies that there are at most $|E_1| + |E_2| - 2 = |E| - 2$ breakpoints and thus at most $|E| - 1$ pieces in the optimal cost function for G .

Finally, consider the case that G is the parallel composition of two series-parallel graphs $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$. Then the optimal cost curve for G is the minimum of the optimal cost curves for G_1 and G_2 . Therefore, the number of its pieces is bounded by the sum of the number of pieces of the curves for G_1 and G_2 , which by induction hypothesis is at most $|E_1| + |E_2| = |E|$. \square

Corollary 2. *There is an algorithm that solves FMP_N for series-parallel graphs in time $O(|E|^2)$.*

Proof. By Theorem 3, the complete optimal cost curve for the parametric shortest path problem with cost function g_λ can be constructed by $O(|E|)$ shortest path computations. Using the linear time shortest path algorithm for planar graphs by Henzinger et al. (1997), the curve can be constructed in time $O(|E|^2)$. \square

3.2. Adaptive Followers' Minimization Problem

The adaptive version of the followers' minimization problem can be solved in polynomial time using a label-setting algorithm. The algorithm makes use of the following observation on the structure of the cost function f_A , with $P[v, w]$ denoting the v - w -subpath of P .

Lemma 3. $f_A(P) = f_A(P[s, v]) + \pi(P[s, v]) \cdot f_A(P[v, t])$.

Proof. Let $P = (e_1, \dots, e_k)$ with $e_i = (v_i, v_{i+1})$ and assume $v = v_{k'}$ for some k' with $1 \leq k' \leq k + 1$. Observe that

$$\begin{aligned} f_A(P) &= \sum_{i=1}^k \prod_{j=1}^{i-1} (1 - p_{e_j}) \cdot (c_{e_i} + p_{e_i}(F + \text{SP}_c(v_{i+1}, t))) \\ &= \sum_{i=1}^{k'-1} \prod_{j=1}^{i-1} (1 - p_{e_j}) \cdot (c_{e_i} + p_{e_i}(F + \text{SP}_c(v_{i+1}, t))) \\ &\quad + \prod_{j=1}^{k'-1} (1 - p_{e_j}) \sum_{i=k'}^k \prod_{j=k'}^{i-1} (1 - p_{e_j}) \\ &\quad \cdot (c_{e_i} + p_{e_i}(F + \text{SP}_c(v_{i+1}, t))) \\ &= f_A(P[s, v]) + \prod_{e \in P[s, v]} (1 - p_e) \cdot f_A(P[v, t]). \quad \square \end{aligned}$$

As an immediate consequence of Lemma 3, we can deduce that every suffix of an optimal path must also be an optimal path. Note, however, that the same is not true for prefixes of the optimal path, as the arcs of $P[s, v]$ also appear in the second summand.

Corollary 3. *Let P^* be an s - t -path minimizing f_A and let $v \in V(P^*)$. Then $P^*[v, t]$ is a v - t -path minimizing f_A .*

In the spirit of Dijkstra's (1959) shortest path algorithm, our algorithm iteratively computes the cost of an optimal v - t -path for some vertex v . It maintains the set S of vertices for which an optimal path has been computed and a value $\phi(v)$ for every vertex $v \in V$, denoting the cost of the cheapest v - t -path found so far. In every iteration, a vertex $w \in V \setminus S$ with minimum value $\phi(w)$ is added to S , and the labels $\phi(v)$ of vertices v with $(v, w) \in E$ are updated if a v - t -path consisting of (v, w) and an optimal w - t -path is cheaper than the current value of $\phi(v)$. A complete listing is given as Algorithm 2.

Algorithm 2 (Algorithm for FMP_A)

```

 $\phi(t) \leftarrow 0$ 
 $\phi(v) \leftarrow \infty$  for all  $v \in V \setminus \{t\}$ 
 $S \leftarrow \emptyset$ 
while  $s \notin S$  do
    Choose  $w \in V \setminus S$  minimizing  $\phi$ .
     $S \leftarrow S \cup \{w\}$ 
    for all  $e = (v, w) \in E_{in}(w)$  do
         $\phi' \leftarrow c_e + p_e(\text{SP}_c(w, t) + F) + (1 - p_e)\phi(w)$ 
        if  $\phi' < \phi(v)$  then
             $\phi(v) \leftarrow \phi'$ 
             $\text{next}(v) \leftarrow e$ 
 $w \leftarrow s; P \leftarrow \emptyset$ 
while  $w \neq t$  do
     $e \leftarrow \text{next}(w); w \leftarrow \text{head}(e)$ 
    Add  $e$  to  $P$ .
return  $P$ 

```

Lemma 4. *When vertex $v \in V$ is added to the set S in Algorithm 2, then there is an optimal v - t -path P_v with $f_A(P_v) = \phi(v)$ starting with the edge $\text{next}(v)$.*

Proof. By contradiction assume the lemma is not true. Without loss of generality, let v be the first vertex added to S that does not fulfill the statement of the lemma and consider the moment when v is added to S . Note that $v \neq t$, as t fulfills the requirements with $\phi(t) = 0$ and $P_t = \emptyset$. Let $\text{next}(v) = e = (v, w)$ for some $w \in S$. As v is the first vertex violating the lemma, there is a w - t -path P_w with $f_A(P_w) = \phi(w)$. Thus, $P_v := (e) \circ P_w$ is a v - t -path with cost

$$f_A(P_v) = c_e + p_e(\text{SP}_c(w, t) + F) + (1 - p_e)\phi(w) = \phi(v).$$

Let P' be any v - t -path, let $e' = (v', w')$ be the last edge on P' with $v' \in V \setminus S$ and $w' \in S$. Note that

$$\begin{aligned} f_A(P'[v', t]) &\geq c_{e'} + p_{e'}(\text{SP}_c(w', t) + F) + (1 - p_{e'})\phi(w') \\ &\geq \phi(v') \geq \phi(v) \end{aligned}$$

where the first inequality follows from the fact that $\phi(w')$ denotes the cost of an optimal w' - t -path by induction hypothesis, the second inequality follows from the fact that $\phi(v')$ was updated when w was added to S , and the last inequality follows from the choice of v . In particular, this implies $\text{SP}_c(v, t) + F \geq f_A(P'[v', t]) \geq \phi(v)$ when choosing P' as a shortest v - t -path with respect to c . Thus, for any arbitrary P' again,

$$\begin{aligned} f_A(P') &\geq (1 - \pi(P'[v, v'])) \cdot (\text{SP}_c(v, t) + F) \\ &\quad + \pi(P'[v, v']) \cdot f_A(P'[v', t]) \geq \phi(v), \end{aligned}$$

which proves that P_v is optimal. \square

Theorem 4. *Algorithm 2 solves FMP_A in time $O(|E| + |V| \log |V|)$.*

Proof. The optimality of the path computed by Algorithm 2 follows immediately from Lemma 4. Shortest path distances from every vertex to t can be pre-computed using Dijkstra's algorithm in time $O(|E| + |V| \log |V|)$. The remainder of the algorithm can be implemented using a Fibonacci heap for computing the vertex minimizing $\phi(v)$, guaranteeing the claimed running time. \square

A Linear Programming Formulation. We conclude our discussion of the adaptive followers' minimization problem by presenting a pair of primal and dual linear programs that describe FMP_A . This formulation will prove useful for constructing exact mixed integer nonlinear programming formulations of the leader's problem in Section 4.3. Let $A_e := c_e + p_e(\text{SP}_c(w, t) + F)$ denote the expected cost of being caught on an edge $e = (v, w) \in E$ and consequently traversing the shortest

path to t . Consider the following pair of primal and dual linear programs.

$$\begin{aligned} [P\text{-FMP}_A] \quad &\min \sum_{e \in E} A_e x_e \\ &\text{s.t.} \quad \sum_{e \in E_{\text{out}}(v)} x_e - \sum_{e \in E_{\text{in}}(v)} (1 - p_e) x_e \\ &\quad = \begin{cases} 1, & \text{if } v = s, \\ 0, & \text{otherwise,} \end{cases} \quad \forall v \in V \setminus \{t\}, \\ &\quad x_e \geq 0, \quad \forall e \in E; \\ [D\text{-FMP}_A] \quad &\max \phi_s \\ &\text{s.t.} \quad \phi_v - (1 - p_e)\phi_w \leq A_e, \quad \forall e = (v, w) \in E; \\ &\quad \phi_t = 0. \end{aligned}$$

These LPs exhibit strong similarities to the primal and dual of the shortest path problem. In fact, $[P\text{-FMP}_A]$ corresponds to a *generalized flow* version of the standard shortest path problem, where x_e denotes the flow on edge $e \in E$ and $(1 - p_e)$ denotes the loss factor; see the thesis of Wayne (1999) for an introduction to generalized flows.

In the same way as standard network flows can be decomposed into flows on paths and cycles, generalized flows can be decomposed into several types of elementary flows. An *elementary flow* $f = (x, P, C)$ consists of a vector $x \in \mathbb{R}_+^E$, a path P , and a cycle C (either of which could be empty), such that $x_e > 0$ only if $e \in P \cup C$. We will write x^f, P^f , and C^f for the vector, path, and cycle associated with elementary flow f . We distinguish three types of elementary flows.

- *Type 1 (path):* An elementary flow f of type 1 is associated with an s - t -path $P^f = (e_1, \dots, e_k) \in \mathcal{P}$, while $C^f = \emptyset$. Its flow vector x^f is defined by $x_{e_1}^f = 1$ and $x_{e_i}^f = (1 - p_{e_{i-1}})x_{e_{i-1}}^f$ for $i \in \{2, \dots, k\}$.

- *Type 2 (path-cycle):* An elementary flow f of type 2 is associated with an s - v -path $P^f = (e_1, \dots, e_l)$ for some $v \in V$ and a cycle $C^f = (e'_1, \dots, e'_l)$ with $e'_i = (v_i, v_{i+1})$ for $i \in \{1, \dots, l\}$ such that $v_1 = v_{l+1} = v$, P^f and C^f are edge disjoint, and $\pi(C^f) := \prod_{i=1}^l (1 - p_{e'_i}) < 1$. Its flow vector x^f is defined by $x_{e_1}^f = 1$, and $x_{e_i}^f = (1 - p_{e_{i-1}})x_{e_{i-1}}^f$ for $i \in \{2, \dots, k\}$, $x_{e'_1}^f = \pi(C^f)/(1 - \pi(C^f))$, and $x_{e'_i}^f = (1 - p_{e'_{i-1}})x_{e'_{i-1}}^f$ for $i \in \{2, \dots, l\}$.

- *Type 3 (loss-less cycle):* An elementary flow f of type 3 is associated with a cycle C^f such that $p_e = 0$ for all $e \in C^f$, while $P^f = \emptyset$. Its flow vector x^f is defined by $x_e^f = 1$ for all $e \in C^f$.

We denote by \mathcal{F}_i the set of elementary flows of type i , and let $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$. It is easy to see that every elementary flow of type 1 or 2 corresponds to a feasible solution to $[P\text{-FMP}_A]$.

Lemma 5. *Let $f \in \mathcal{F}_1 \cup \mathcal{F}_2$. Then x^f is a feasible solution to $[P\text{-FMP}_A]$.*

Proof. For $w \in V$, define $\text{ex}(w) := \sum_{e \in E_{\text{out}}(w)} x_e^f - \sum_{e \in E_{\text{in}}(w)} (1 - p_e) x_e^f$. Note that $\text{ex}(s) = x_{e_1}^f = 1$, regardless of the type of f . It remains to show that $\text{ex}(w) = 0$ for all $w \in V \setminus \{s, t\}$. If f is of type 1, or if f is of type 2 and $v \neq w$, this is immediately obvious by construction of f . If f is of type 2 and $w = v$, observe that $\text{ex}(v) = x_{e_1}^f - (1 - p_{e_k}) x_{e_k}^f - (1 - p_{e_1}) x_{e_1}^f$, where $(1 - p_{e_k}) x_{e_k}^f = \pi(P^f)$ and $(1 - p_{e_1}) x_{e_1}^f = \pi(C) x_{e_1}^f$ by construction of f . Inserting the definition of $x_{e_1}^f$ yields $\text{ex}(v) = 0$, and hence x^f is a feasible solution to $[P\text{-FMP}_A]$. \square

Conversely, any feasible solution to $[P\text{-FMP}_A]$ can be written as a convex combination of elementary flows of types 1 and 2, plus possibly a conic combination of elementary flows of type 3. A decomposition of a generalized flow x is a vector $\mu^f \in \mathbb{R}_+^{\mathcal{F}}$ such that $x_e = \sum_{f \in \mathcal{F}} \mu^f x_e^f$. Note that $\sum_{f \in \mathcal{F}_1 \cup \mathcal{F}_2} \mu^f$ corresponds to the net outflow of the source s . The following lemma is an immediate consequence of the classic decomposition theorem for generalized flows (Gondran et al. 1984).

Lemma 6. *Let x be a feasible solution to $[P\text{-FMP}_A]$. Then there is a decomposition μ of x with $\sum_{f \in \mathcal{F}_1 \cup \mathcal{F}_2} \mu^f = 1$.*

Proof. We apply the decomposition theorem for generalized flows in the version given as in the thesis of (Wayne 1999, Theorem 2.3.3). Note that, because $1 - p_e \leq 1$ for all $e \in E$, there are no flow-generating cycles in the network. Hence, the theorem decomposes the generalized flow into elementary flows of types I, II, and IV in the notation of Wayne (1999), corresponding to types 1, 3, and 2 in our notation, respectively. Although the flow values on the arcs of the elementary flows are not explicitly stated by Wayne (1999), it is easy to see that the values given in our definition above are the unique values fulfilling flow conservation w.r.t. the gain factors $1 - p_e$ when normalizing the flow value of the first arc to 1. As every elementary flow of type 1 or 2 creates a deficit of 1 at the source s , and also x has a deficit of 1 at s , we conclude that $\sum_{f \in \mathcal{F}_1 \cup \mathcal{F}_2} \mu^f = 1$. \square

Furthermore, we can also relate the cost of any elementary flow to $f_A(P)$ for some s - t -path $P \in \mathcal{P}$.

Lemma 7. *If $f \in \mathcal{F}_1$, then $f_A(P^f) = \sum_{e \in P^f} A_e x_e^f$. If $f \in \mathcal{F}_2$, then there is an s - t -path $P \in \mathcal{P}$ with $f_A(P) \leq \sum_{e \in E} A_e x_e^f$. If the latter holds with equality, then $\pi(P) = 0$.*

Proof. If f is of type 1, by definition of A_e and x_e^f , we see that

$$\begin{aligned} \sum_{e \in P^f} A_e x_e^f &= \sum_{j=1}^k (c_{e_j} + p_{e_j} (\text{SP}_c(v_{j+1}, t) + F)) \prod_{i=1}^{j-1} (1 - p_{e_i}) \\ &= f_A(P^f). \end{aligned}$$

Now suppose that f is of type 2. Recall that f corresponds to an s - v -path $P^f = (e_1, \dots, e_k)$ that meets a

cycle $C^f = (e'_1, \dots, e'_l)$ at a node v . Define $\pi_j := \prod_{i=1}^j (1 - p_{e'_i})$. Using exactly the same transformations as in the proof of Lemma 1, we obtain

$$\begin{aligned} &\sum_{j=1}^l \pi_{j-1} (c_{e'_j} + p_{e'_j} (F + \text{SP}_c(v_{j+1}, t))) \\ &\geq \pi_l \sum_{j=1}^l c_{e'_j} + (1 - \pi_l) (F + \text{SP}_c(v, t)). \end{aligned}$$

Multiplying the above inequality with $\pi(P^f)/(1 - \pi_l)$, using the identity $x_{e'_j} = (\pi(P^f)/(1 - \pi_l)) \pi_{j-1}$, and dropping the first term on the right-hand side yields $\sum_{i=1}^l A_{e'_i} x_{e'_i} \geq \pi(P^f) (F + \text{SP}_c(v, t))$. Now let \bar{P} be a shortest v - t -path with respect to c . Note that $f_A(\bar{P}) = \text{SP}_c(v, t) + (1 - \pi(\bar{P}))F$ by Lemma 1. Consider the concatenation $P := P^f \circ \bar{P}$. Using Lemma 3, we obtain

$$\begin{aligned} f_A(P) &= f_A(P^f) + \pi(P^f) f_A(\bar{P}) \\ &= \sum_{i=1}^k A_{e_i} \underbrace{\prod_{j=1}^{i-1} (1 - p_{e_j})}_{=x_{e_i}} + \pi(P^f) \cdot (\text{SP}_c(v, t) + (1 - \pi(\bar{P}))F) \\ &= \sum_{e \in P^f} A_e x_e + \pi(P^f) (F + \text{SP}_c(v, t)) - \pi(P^f) \pi(\bar{P}) F \\ &\leq \sum_{e \in P^f} A_e x_e + \sum_{e \in C^f} A_e x_e - \pi(P) F. \quad \square \end{aligned}$$

Theorem 5. *For every s - t -path P , there is a feasible solution x to $[P\text{-FMP}_A]$ with $f_A(P) = \sum_{e \in E} A_e x_e$ and $\pi(P) = \sum_{e \in E_{\text{in}}(t)} (1 - p_e) x_e$. For every optimal solution x to $[P\text{-FMP}_A]$, there is an s - t -path P with $f_A(P) = \sum_{e \in E} A_e x_e$ and $\pi(P) \leq \sum_{e \in E_{\text{in}}(t)} (1 - p_e) x_e$.*

Proof. The first statement of the theorem follows immediately from Lemma 7, using the elementary flow f of type 1 with $P^f = P$. The second part follows by applying Lemma 6 to obtain a decomposition μ of x and noting that $\sum_{e \in E} A_e x_e = \sum_{f \in \mathcal{F}} \mu^f \sum_{e \in E} A_e x_e^f$ with $\mu^f \geq 0$ for all $f \in \mathcal{F}$ and $\sum_{f \in \mathcal{F}_1 \cup \mathcal{F}_2} \mu^f = 1$ implies that $\sum_{e \in E} A_e x_e = \sum_{e \in E} A_e x_e^f$ for all $f \in \mathcal{F}_1 \cup \mathcal{F}_2$ with $\mu^f > 0$. Choose $f \in \mathcal{F}_1 \cup \mathcal{F}_2$ such that $\sum_{e \in E_{\text{in}}(t)} (1 - p_e) x_e^f$ is minimal. If f is of type 1, then $f_A(P^f) = \sum_{e \in E} A_e x_e^f$ by Lemma 7, and $\pi(P^f) = \prod_{i=1}^k (1 - p_{e_i}) = (1 - p_{e_k}) x_{e_k}^f \leq \sum_{e \in E_{\text{in}}(t)} (1 - p_e) x_e$, where the last inequality is due to the minimality assumption and the fact that e_k is the only incoming arc of t in the support of x^f . If f is of type 2, then, again by Lemma 7 there is a path $P \in \mathcal{P}$ with $f_A(P) \leq \sum_{e \in E} A_e x_e^f$. As, by optimality of x , this inequality must actually hold with equality, and thus Lemma 7 guarantees $\pi(P) = 0 \leq \sum_{e \in E_{\text{in}}(t)} (1 - p_e) x_e$. \square

3.3. The Impact of Adaptivity

We now prove a tight upper bound of $4/3$ on the ratio of the optimal cost between nonadaptive and adaptive

strategies. In fact, we prove a slightly stronger bound, stated in the following lemma, which we will also use in later sections.

Lemma 8. Let OPT_N be the cost of an optimal nonadaptive solution and OPT_A be the cost of an optimal adaptive solution. Then $\text{OPT}_N - \text{SP}_c(s, t) \leq \frac{4}{3}(\text{OPT}_A - \text{SP}_c(s, t))$.

Proof. Let P be a path that minimizes f_A . Observe that

$$\text{OPT}_A = f_A(P) \geq \pi(P)c(P) + (1 - \pi(P))(\text{SP}_c(s, t) + F) \quad (1)$$

by Lemma 1. Furthermore,

$$\text{OPT}_N \leq \min\{c(P) + (1 - \pi(P))F, \text{SP}_c(s, t) + F\}, \quad (2)$$

as both P and a shortest path from s to t are feasible nonadaptive solutions. Note that we can assume $\pi(P) < 1$ and $F > 0$ as otherwise $\text{OPT}_A = \text{OPT}_N$. Thus, introducing variables C and S and x corresponding to $c(P)$, $\text{SP}_c(s, t)$, and $1 - \pi(P)$, respectively, in Equations (1) and (2), we obtain

$$\begin{aligned} \frac{\text{OPT}_N - \text{SP}_c(s, t)}{\text{OPT}_A - \text{SP}_c(s, t)} &\leq \max_{\substack{x \in (0, 1], F > 0 \\ 0 \leq S \leq C}} \frac{\min\{C + xF, S + F\} - S}{(1 - x)C + x(S + F) - S} \\ &= \max_{\substack{x \in (0, 1], F > 0 \\ 0 \leq S \leq C}} \frac{\min\{C - S + xF, F\}}{(1 - x)(C - S) + xF}, \end{aligned}$$

where the first inequality follows from the fact that setting $C = c(P)$, $S = \text{SP}_c(s, t)$, and $x = 1 - \pi(P)$ is a feasible solution to the maximization problem. We now argue that, in an optimal solution to the maximization problem, $C - S + xF = F$. If $C - S + xF > F$, then increasing S also increases the value of the right-hand side function. If $C - S + xF < F$, then decreasing F increases the value of the right-hand side function. Thus, if the right-hand side is maximized, then $C - S + xF = F$. Therefore,

$$\frac{\text{OPT}_N - \text{SP}_c(s, t)}{\text{OPT}_A - \text{SP}_c(s, t)} \leq \max_{\substack{x \in (0, 1] \\ F > 0}} \frac{F}{F(1 - x + x^2)} \leq \frac{4}{3},$$

which concludes the proof. \square

Theorem 6. Let OPT_N be the cost of an optimal nonadaptive solution and OPT_A be the cost of an optimal adaptive solution. Then $\text{OPT}_N \leq \frac{4}{3}\text{OPT}_A$.

Proof. This follows immediately from Lemma 8 by observing that if $\text{OPT}_A = \text{SP}_c(s, t)$ then also $\text{OPT}_N = \text{SP}_c(s, t)$ and that otherwise $\text{OPT}_N/\text{OPT}_A \leq (\text{OPT}_N - \text{SP}_c(s, t))/(\text{OPT}_A - \text{SP}_c(s, t))$. \square

The following example shows that the bound given in Theorem 6 is tight.

Example 1. Let $V = \{s, v, t\}$ and $E = \{e_0, e_1, e_2\}$ where e_0 goes from s to v and e_1 and e_2 are two parallel arcs from v to t . The travel costs are $c_{e_0} = c_{e_1} = 0$ and $c_{e_2} = 1$, the inspection probabilities are $p_{e_0} = 1/2$, $p_{e_1} = 1$, and $p_{e_2} = 0$, and $F = 2$. Observe that there are only two s - t -paths $P_1 = (e_0, e_1)$ and $P_2 = (e_0, e_2)$ and that $f_N(P_1) = f_N(P_2) = 2$ and $f_A(P_2) = \frac{3}{2}$, yielding a ratio of $\frac{4}{3}$ between optimal nonadaptive and adaptive strategies.

Remark 1. Note that in the proof of Theorem 6, we did not make use of the fact that the probability $\pi(P)$ is determined by individual probabilities on the arcs of the network. Therefore, the bound of $\frac{4}{3}$ given by the theorem is still true for arbitrary probability distributions specified by a probability $\pi(P)$ for each path P . In particular, the result still holds if the inspections at the arcs are not independent events.

4. The Leader's Maximization Problem

In this section we discuss algorithms and complexity results for the leader's maximization problem. On the theoretical side, we derive NP -hardness for a restricted special case of the problem and an LP relaxation, which yields upper bounds on the profit for all four model variants. For the flexible-fare setting, we also obtain a constant factor approximation. On the practical side, we propose a local search procedure that, combined with initial solutions from the LP relaxation computes close-to-optimal solutions for all four variants of the problem.

4.1. Complexity of the Leader's Problem

All four variants of the leader's maximization problem are NP -hard even in very restricted cases, as can be seen from a simple reduction from the minimum directed multicut problem.

Theorem 7. LMP_X^L for $L \in \{\text{fix}, \text{flex}\}$ and $X \in \{A, N\}$ is strongly NP -hard, even when restricted to instances with $|K| = 2$ and $c \equiv 0$.

Proof. Consider an instance $I = (G, (s_i, t_i)_{i=1, \dots, k}, q)$ of the directed multicut problem with $G = (V, E)$, $k = 2$, and $q \in \mathbb{Z}_+$. We will construct an instance $\hat{I} = (G, c, K, F, T, B)$ of the leader's maximization problem as follows. We introduce two commodities, one for each pair (s_i, t_i) with $i \in K := \{1, 2\}$. We set $c_e = 0$ for all $e \in E$, $B = q$ and $T_1 = T_2 = F = 1$. We denote by OPT_X^L the value of an optimal solution to LMP_X^L . Note that since the travel costs are all zero, $f_{N,p,i} = f_{A,p,i}$ for all $i \in K$ and any setting of probabilities $p \in [0, 1]^E$. Furthermore, $\Gamma_{N,i}^{\text{fix}}(p) = \min_{P \in \mathcal{P}_i} \{f_{N,p,i}(P), T_i\} = \Gamma_{N,i}^{\text{flex}}(p)$ as $\text{SP}_c(s_i, t_i) = 0$ for all $i \in K$. Therefore, $\text{OPT}_N^{\text{fix}} = \text{OPT}_N^{\text{flex}} = \text{OPT}_N^{\text{flex}} = \text{OPT}_A^{\text{flex}}$. We show that $\text{OPT}_X^L \geq 2$ if and only if there exists a feasible multicut of cardinality q , proving the theorem.

Suppose that there is a multicut M with $|M| \leq q$. We then define $p_e = 1$ for all $e \in M$ and $p_e = 0$ for all $e \in E \setminus M$. Note that $\sum_{e \in E} p_e = |M| \leq q$, implying that p is a feasible solution. Furthermore, every passenger encounters an inspector with probability 1 because his path has to cross the multicut. Thus, we obtain $\text{OPT}_X^L \geq 2$.

Conversely, assume $p \in [0, 1]^E$ is a solution with profit 2. This implies that $f_{X,i,p}(P) = 1$ for $i \in K$ and every path $P \in \mathcal{P}_i$. This is only possible if, for every $P \in \mathcal{P}_i$, there is an $e \in P$ with $p_e = 1$. Therefore, the set $M = \{e \in E : p_e = 1\}$ is a multicut with cardinality $|M| \leq \sum_{e \in M} p_e \leq q$. \square

4.2. LP Relaxation and Approximation

Let OPT_X^L denote the value of an optimal solution to the corresponding version of the leader's maximization problem for $L \in \{\text{flex}, \text{fix}\}$ and $X \in \{A, N\}$. The following lemmas relate these values to one another.

Lemma 9. $\text{OPT}_N^{\text{flex}} \geq \text{OPT}_N^{\text{fix}}$ and $\text{OPT}_A^{\text{flex}} \geq \text{OPT}_A^{\text{fix}}$.

Proof. Let $X \in \{A, N\}$. Note that $f_{X,p,i}(P) - \text{SP}_c(s_i, t_i) \geq (1 - \pi(P))F$ for all $p \in [0, 1]^E$, $i \in K$ and any path $P \in \mathcal{P}_i$ by Corollary 1. Therefore $\text{OPT}_X^{\text{flex}} \geq \text{OPT}_X^{\text{fix}}$. \square

Lemma 10. $\frac{4}{3}\text{OPT}_A^{\text{flex}} \geq \text{OPT}_N^{\text{flex}} \geq \text{OPT}_A^{\text{flex}}$.

Proof. Let $p \in [0, 1]^E$ be any solution to the LMP instance. Note that Lemma 8 implies $\frac{4}{3}\Gamma_{A,i}^{\text{flex}}(p) \geq \Gamma_{N,i}^{\text{flex}}(p) \geq \Gamma_{A,i}^{\text{flex}}(p)$ for all $i \in K$, and therefore $\frac{4}{3}\text{OPT}_A^{\text{flex}} \geq \text{OPT}_N^{\text{flex}} \geq \text{OPT}_A^{\text{flex}}$. \square

To obtain an LP relaxation, we will make use of a linearization approach, which is based on the following classic approximation.

Lemma 11. $1 - \pi(P) \leq \min\{\sum_{e \in P} p_e, 1\} \leq 1/(1 - e^{-1}) \cdot (1 - \pi(P))$.

Proof. We prove $\pi(P) = \prod_{e \in P} (1 - p_e) \geq 1 - \sum_{e \in P} p_e$ by induction on $|P|$. This is trivial for $|P| = 1$. For $|P| > 1$, observe that for any $e' \in P$ by induction hypothesis

$$\begin{aligned} \prod_{e \in P} (1 - p_e) &\geq \left(1 - \sum_{e \in P \setminus \{e'\}} p_e\right) (1 - p_{e'}) \\ &= 1 - \sum_{e \in P} p_e + p_{e'} \sum_{e \in P \setminus \{e'\}} p_e \geq 1 - \sum_{e \in P} p_e. \end{aligned}$$

This immediately implies the first inequality stated in the lemma.

The second inequality of the lemma is trivially true if $\sum_{e \in P} p_e = 0$. Thus, we assume the sum to be strictly positive without loss of generality. Define $\sigma := \min\{\sum_{e \in P} p_e, 1\}$ and observe that

$$\prod_{e \in P} (1 - p_e) \leq \left(1 - \frac{\sigma}{|P|}\right)^{|P|} \leq e^{-\sigma}.$$

Therefore,

$$\frac{1 - \pi(P)}{\sigma} \geq \frac{1 - e^{-\sigma}}{\sigma} \geq \min_{x \in (0, 1]} \frac{1 - e^{-x}}{x}.$$

The right-hand side is decreasing in x and therefore minimized for $x = 1$. \square

Using Lemma 11, we replace the term $1 - \pi(P)$ in the followers' objective function by $\sum_{e \in P} p_e$. Note that after this replacement, the nonadaptive version of FMP corresponds to a classic shortest path problem. Using the

dual of the shortest path LP, we derive the following LP relaxation for $\text{LMP}_N^{\text{flex}}$.

$$\begin{aligned} \text{[LP]} \quad \max \quad & \sum_{i \in K} d_i(y_i(t_i) - y_i(s_i) - \text{SP}_c(s_i, t_i)) \\ \text{s.t.} \quad & \sum_{e \in E} p_e \leq B; \\ & y_i(w) - y_i(v) \leq c_e + Fp_e, \quad \forall i \in K, e = (v, w) \in E; \\ & y_i(v) - y_i(s_i) \leq F + \text{SP}_c(s_i, v), \quad \forall i \in K, v \in V; \\ & p_e \in [0, 1], \quad \forall e \in E. \end{aligned}$$

The value OPT_{LP} of an optimal solution to [LP] yields an upper bound to all four variants of the leader's maximization problem.

Lemma 12. $\text{OPT}_{\text{LP}} \geq \text{OPT}_X^L$ for all $L \in \{\text{flex}, \text{fix}\}$ and $X \in \{A, N\}$.

Proof. Let $p \in [0, 1]^E$ be an optimal solution to $\text{LMP}_N^{\text{flex}}$. For every $i \in K$, set $y_i(s_i) = 0$ and $y_i(v) = \min\{\text{SP}_{c+Fp}(s_i, v), \text{SP}_c(s_i, v) + F\}$ for all $v \in V \setminus \{s_i\}$. It is easy to check that (p, y) is a feasible solution to [LP] and that $y_i(t_i) - y_i(s_i) \geq \min\{\text{SP}_{c+Fp}(s_i, t_i), F + \text{SP}_c(s_i, t_i)\}$ for every $i \in K$. Therefore,

$$\begin{aligned} & y_i(t_i) - y_i(s_i) - \text{SP}_c(s_i, t_i) \\ & \geq \min\left\{\sum_{e \in P} (c_e + Fp_e), F + \text{SP}_c(s_i, t_i) : P \in \mathcal{P}_i\right\} \\ & \quad - \text{SP}_c(s_i, t_i) \\ & \geq \min\left\{\sum_{e \in P} c_e + (1 - \pi(P))F : P \in \mathcal{P}_i\right\} - \text{SP}_c(s_i, t_i) \\ & = \min\{f_{N,i,p}(P) - \text{SP}_c(s_i, t_i) : P \in \mathcal{P}_i\} = \Gamma_{N,i}^{\text{flex}}(p), \end{aligned}$$

where the second inequality follows from Lemma 11 and the equality from the observation that $\min\{f_{N,i,p}(P) : P \in \mathcal{P}_i\} \leq \text{SP}_c(s_i, t_i) + F$. This implies that the optimal value of the LP is at least $\text{OPT}_N^{\text{flex}}$. By Lemmas 9 and 10, $\text{OPT}_N^{\text{flex}}$ is as least as large as the optimal solution value of any of the other versions. \square

Using Lemmas 10 and 11, we can also derive that using an optimal solution to [LP] yields approximation algorithms for $\text{LMP}_N^{\text{flex}}$ and $\text{LMP}_A^{\text{flex}}$.

Theorem 8. *There is a $(1 - 1/e)$ -approximation algorithm for $\text{LMP}_N^{\text{flex}}$.*

Proof. Let (p, y) be an optimal solution to [LP]. Note that $y_i(t_i) - y_i(s_i) \leq \min\{\text{SP}_{c+Fp}(s_i, t_i), F + \text{SP}_c(s_i, t_i)\}$ and define $\lambda_i := y_i(t_i) - y_i(s_i) - \text{SP}_c(s_i, t_i)$. Then

$$\begin{aligned} \lambda_i &\leq \min\left\{\sum_{e \in P} (c_e + Fp_e) - \text{SP}_c(s_i, t_i), F\right\} \\ &\leq \sum_{e \in P} c_e - \text{SP}_c(s_i, t_i) + \min\left\{\sum_{e \in P} p_e, 1\right\} \cdot F \\ &\leq \frac{1}{1 - e^{-1}} \left(\sum_{e \in P} c_e - \text{SP}_c(s_i, t_i) + \left(1 - \prod_{e \in P} (1 - p_e)\right) \cdot F\right) \end{aligned}$$

for every $P \in \mathcal{P}_i$ by Lemma 11. Thus, setting the probabilities according to p yields a solution to $\text{LMP}_N^{\text{flex}}$ with profit

$$\sum_{i \in K} d_i \Gamma_{N,i}^{\text{flex}}(p) \geq (1 - 1/e) \sum_{i \in K} d_i \lambda_i \geq (1 - 1/e) \text{OPT}_N^{\text{flex}}. \quad \square$$

By Lemma 8 this result also translates to $\text{LMP}_A^{\text{flex}}$, with a loss of a factor of 3/4 in the approximation guarantee.

Corollary 4. *There is a $\frac{3}{4}(1 - 1/e)$ -approximation algorithm for $\text{LMP}_A^{\text{flex}}$.*

Remark 2. The analysis of the algorithm given in Theorem 8 is tight. To see this, consider the following example instance of $\text{LMP}_N^{\text{flex}}$. Let G be a directed cycle of length n , i.e., $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_n\}$ with $e_i = (v_i, v_{i+1})$ for $i \in \{1, \dots, n-1\}$ and $e_n = (v_n, v_1)$. Let K consist of n commodities with unit demand, such that $s_i = v_i$ and $t_i = v_{i-1}$ for $i \in \{2, \dots, n\}$ and $s_1 = v_1$ and $t_1 = v_n$. Note \mathcal{P}_i consists of a unique path of length $n-1$ for every $i \in K$. Finally, let $c = 0$, $F = T_i = 1$ for every $i \in K$ and $B = n/(n-1)$. Observe that the optimal solution of [LP] sets $p_e = 1/(n-1)$ for every $e \in E$. Using these probabilities as a solution to $\text{LMP}_N^{\text{flex}}$ yields a profit of $n \cdot (1 - (1 - 1/(n-1))^{n-1})$. On the other hand, setting $p_{e_1} = 1$ and $p_{e_n} = 1/(n-1)$ yields a profit of $n-1 + 1/(n-1)$. Note that by choosing n sufficiently large the ratio between these two values can be brought arbitrarily close to $1 - 1/e$.

Remark 3. Unfortunately, [LP] does not yield an approximation guarantee for the fixed-fare setting. To see this, consider the following instance of $\text{LMP}_X^{\text{fix}}$. There are four nodes s_1, t_1, s_2, t_2 , together with edges from s_1 to s_2 , s_2 to t_2 , and t_2 to t_1 , each with zero cost. In addition, there are L parallel edges from s_1 to t_1 , each with cost 1. Commodity 1 has origin s_1 and destination t_1 with a total demand of L . Commodity 2 has origin s_2 and destination t_2 with a total demand of 1. The budget is $1/2 + \varepsilon$, and the fine is 2. Observe that the optimal solution of the corresponding instance of [LP] sets $p_{(s_2, t_2)} = 1/2 + \varepsilon$ and $p_e = 0$ for all other edges. Interpreting these probabilities as a solution of $\text{LMP}_X^{\text{fix}}$ results in a profit of $1 + 2\varepsilon$, as the followers represented by commodity 1 will prefer one of the edges from s_1 to t_1 over the three-edge path. However, setting $p_{(s_2, t_2)} = 1/2$ instead yields a profit of $L + 1$.

4.3. Exact Nonlinear Formulations

Based on the LP formulation of FMP_A , we can construct exact nonlinear formulations for LMP_A^L for all $L \in \{\text{flex}, \text{fix}\}$. Furthermore, using a path formulation to model the followers' problem, we can construct exact

mixed integer nonlinear formulations for LMP_N^L for all $L \in \{\text{flex}, \text{fix}\}$.

4.3.1. A Nonlinear Formulation for $\text{LMP}_A^{\text{flex}}$. In Section 8 we presented a compact linear program, $[P\text{-FMP}_A]$, and its dual, $[D\text{-FMP}_A]$, describing the adaptive follower's objective. Observe that for the flexible ticket prices setting, the value of the leader's objective function only depends on $f_{X,p,i}$, and not the actual routing decisions of the followers. Therefore, it suffices to turn the probabilities in the formulation into decision variables and maximize over the difference of the node potentials, $\phi_i(v)$, for each node $v \in V$ and commodity $i \in K$, subject to the budget constraint. Applying this transformation to the linear program yields the following nonlinear, compact formulation for $\text{LMP}_A^{\text{flex}}$. Let $A_{ie}(p) := c_e + p_e(\text{SP}_c(w, t_i) + F)$ denote the expected cost of being caught on an edge $e = (v, w) \in E$ and consequently traversing the shortest path to t_i .

$$\begin{aligned} [\text{NLP}_A^{\text{flex}}] \quad & \max \sum_{i \in K} d_i (\phi_i(s_i) - \text{SP}_c(s_i, t_i)) \\ & \text{s.t. } \phi_i(v) - (1 - p_e) \phi_i(w) \leq A_{ie}(p) \\ & \quad \quad \quad \forall i \in K, e = (v, w) \in E \\ & \phi_i(t_i) = 0 \quad \forall i \in K \\ & \sum_{e \in E} p_e \leq B \\ & p_e \in [0, 1], \quad \forall e \in E. \end{aligned}$$

Lemma 13. *Let $\text{OPT}_{[\text{NLP}_A^{\text{flex}}]}$ denote the value of an optimal solution to $[\text{NLP}_A^{\text{flex}}]$. Then $\text{OPT}_{[\text{NLP}_A^{\text{flex}}]} = \text{OPT}_A^{\text{flex}}$.*

Proof. Let $p \in [0, 1]^E$ with $\sum_{e \in E} p_e \leq B$. Let ϕ be such that it maximizes the objective of $[\text{NLP}_A^{\text{flex}}]$ when keeping p fixed. Note that for $i, j \in K$ with $i \neq j$, the values of ϕ_i and ϕ_j can be optimized independently. Thus, for each $i \in K$ the value of ϕ_i is an optimal solution to an instance of $[D\text{-FMP}_A]$ for the given value of p . Therefore, by Theorem 5 and LP duality, $\phi_i(s_i) = f_{A,p,i}(P_i)$, where P_i is an optimal s_i - t_i -path for follower i given probabilities p . Hence, the LP objective value of the solution (p, ϕ) is

$$\begin{aligned} \sum_{i \in K} d_i (\phi_i(s_i) - \text{SP}_c(s_i, t_i)) &= \sum_{i \in K} d_i (f_{A,p,i}(P_i) - \text{SP}_c(s_i, t_i)) \\ &= \sum_{i \in K} d_i \Gamma_{A,i}^{\text{flex}}(p). \end{aligned}$$

Therefore, any optimal solution to $[\text{NLP}_A^{\text{flex}}]$ corresponds to an optimal solution of $\text{LMP}_A^{\text{flex}}$. \square

4.3.2. A Nonlinear Formulation for $\text{LMP}_A^{\text{fix}}$. Since for $\text{LMP}_A^{\text{fix}}$, the leader maximizes over the revenue gained from expected fines (which depends on the paths traversed by the commodities), the above transformation to a nonlinear formulation for the leader's problem does not work for the fixed setting. However, using

the decomposition lemma from Section 8, we can write down a single level nonlinear programming formulation for $\text{LMP}_A^{\text{fix}}$, following the general idea of Labbé et al. (1998). Taking the natural bilevel mathematical program for the problem, we replace the lower level program by the constraints appearing in $[D\text{-FMP}_A]$ and $[P\text{-FMP}_A]$, and a constraint ensuring that the objective function of the primal is equal to that of the dual. Furthermore, the expected fines in the objective function are calculated based on the flow on the paths from the generalized flow decomposition.

$$\begin{aligned}
[\text{NLP}_A^{\text{fix}}] \quad & \max \sum_{i \in K} d_i \left(1 - \sum_{e \in E_{\text{in}}(t_i)} (1 - p_e) x_{ie} \right) F \\
\text{s.t.} \quad & \phi_i(s_i) = \sum_{e=(v,w) \in E} x_{ie} A_{ie}, \quad \forall i \in K; \\
& \phi_i(v) - (1 - p_e) \phi_i(w) \leq A_{ie}, \\
& \quad \quad \quad \forall i \in K, e = (v, w) \in E; \\
& \sum_{e \in E_{\text{out}}(v)} x_{ie} - \sum_{e \in E_{\text{in}}(v)} (1 - p_e) x_{ie} \\
& = \begin{cases} 1, & \text{if } v = s_i, \\ 0, & \text{otherwise,} \end{cases} \quad \forall i \in K, v \in V; \\
& \sum_{e \in E} p_e \leq B; \\
& \phi_i(t_i) = 0, \quad \forall i \in K; \\
& x_{ie} \geq 0, \quad \forall i \in K, e \in E; \\
& p_e \in [0, 1], \quad \forall e \in E.
\end{aligned}$$

Lemma 14. Let $\text{OPT}_{[\text{NLP}_A^{\text{fix}}]}$ denote the value of an optimal solution to $[\text{NLP}_A^{\text{fix}}]$. Then $\text{OPT}_{[\text{NLP}_A^{\text{fix}}]} = \text{OPT}_A^{\text{fix}}$.

Proof. Let $p \in [0, 1]^E$ with $\sum_{e \in E} p_e \leq B$. Let (x, ϕ) be an optimal solution to $[\text{NLP}_A^{\text{fix}}]$ when keeping p fixed. By LP duality, x_i is an optimal solution to $[P\text{-FMP}_A]$ and ϕ_i is an optimal solution to $[D\text{-FMP}_A]$ for each $i \in K$. Thus, by Theorem 5, for each $i \in K$ there exists an s_i - t_i -path P_i with $f_{A,p,i}(P_i) = \min_{P \in \mathcal{P}_i} f_{A,p,i}(P)$ and $\pi(P_i) \leq \sum_{e \in E_{\text{in}}(t_i)} (1 - p_e) x_{ie}$. Since the objective is maximized, $\sum_{e \in E_{\text{in}}(t_i)} (1 - p_e) x_{ie}$ is minimum among all optimal solutions to $[P\text{-FMP}_A]$ for follower i . The first part of Theorem 5, therefore, guarantees that $\pi(P_i) \geq \sum_{e \in E_{\text{in}}(t_i)} (1 - p_e) x_{ie}$. Hence, we obtain

$$\begin{aligned}
\sum_{i \in K} d_i \left(1 - \sum_{e \in E_{\text{in}}(t_i)} (1 - p_e) x_{ie} \right) F &= \sum_{i \in K} d_i (1 - \pi(P_i)) F \\
&= \sum_{i \in K} d_i \Gamma_{A,i}^{\text{flex}}(p).
\end{aligned}$$

As $p \in [0, 1]^E$ can be chosen freely within the budget constraints, an optimal solution to $[\text{NLP}_A^{\text{fix}}]$ corresponds to an optimal solution of $\text{LMP}_A^{\text{fix}}$ and vice versa. \square

4.3.3. A Mixed Integer Nonlinear Model for LMP_N^L . We cannot straightforwardly apply the same techniques as above for LMP_N^L , since it is unclear how to formulate optimality conditions for FMP_N . Therefore, we introduce a path-based formulation of LMP_N^L in which there is a binary variable x_P for every path $P \in \mathcal{P}_i$ and every $i \in K$, encoding which path follower i takes in the network. This modeling step allows to replace the lower level program by linear constraints encoding the followers' optimality conditions. To encode the leader's objective (with either fixed or flexible ticket prices), we define $C_i^{\text{flex}}(p, x) := \sum_{P \in \mathcal{P}_i} (f_{N,p,i}(P) - \text{SP}_c(s_i, t_i)) x_P$ and $C_i^{\text{fix}}(p, x) := \sum_{P \in \mathcal{P}_i} (1 - \pi(P)) F x_P$ for every $i \in K$. We also let $\mathcal{P}_i^{\text{fix}} := \mathcal{P}_i \cup \{P_i^T\}$ (with P_i^T as described in Section 2.2) and $\mathcal{P}_i^{\text{flex}} := \mathcal{P}_i$. We obtain the following (single level) mixed integer nonlinear program (MINLP) for $L \in \{\text{fix}, \text{flex}\}$.

$$\begin{aligned}
[\text{MINLP}_N^L] \quad & \max \sum_{i \in K} d_i C_i^L(p, x) \\
\text{s.t.} \quad & \sum_{P' \in \mathcal{P}_i^L} f_{N,p,i}(P') x_{P'} \leq f_{N,p,i}(P), \\
& \quad \quad \quad \forall i \in K, P \in \mathcal{P}_i^L; \\
& \sum_{P \in \mathcal{P}_i^L} x_P = 1, \quad \forall i \in K; \\
& \sum_{e \in E} p_e \leq B; \\
& x_P \in \{0, 1\}, \quad \forall i \in K, P \in \mathcal{P}_i; \\
& p_e \in [0, 1], \quad \forall e \in E.
\end{aligned}$$

Lemma 15. Let $\text{OPT}_{[\text{MINLP}_N^{\text{fix}}]}$ and $\text{OPT}_{[\text{MINLP}_N^{\text{flex}}]}$ denote the value of an optimal solution to $[\text{MINLP}_N^{\text{fix}}]$ and $[\text{MINLP}_N^{\text{flex}}]$, respectively. Then $\text{OPT}_{[\text{MINLP}_N^{\text{fix}}]} = \text{OPT}_N^{\text{fix}}$ and $\text{OPT}_{[\text{MINLP}_N^{\text{flex}}]} = \text{OPT}_N^{\text{flex}}$.

Proof. Let $p \in [0, 1]^E$ with $\sum_{e \in E} p_e \leq B$. Let x be an optimal solution to $[\text{MINLP}_N^L]$ for $L \in \{\text{fix}, \text{flex}\}$ when keeping p fixed. Observe that for each $i \in K$ there is a unique path P_i with $x_{P_i} = 1$. By the first constraint, P_i minimizes the follower i 's objective $f_{N,p,i}$, i.e., $P_i \in \arg \min_{P \in \mathcal{P}_i^L} f_{N,p,i}(P)$. As the objective of $[\text{MINLP}_N^L]$ is maximized, P_i is chosen such that it maximizes the leader's profit among all paths in $\arg \min_{P \in \mathcal{P}_i^L} f_{N,p,i}(P)$. Therefore, $C_i^L(p, x) = \Gamma_{N,i}^L(p)$. As p can be chosen freely within the budget constraints, the statement of the theorem follows. \square

4.4. Local Search Framework

We conclude this section by presenting a general local search framework to compute close-to-optimal solutions to the leader's maximization problem. The approach can be applied to all four model variants

by using the corresponding followers' response and leader's objective function.

Algorithm 3 (Local search).

Input: instance I of LMP_X^L , initial solution p , revenue algorithm REVENUE.
Output: probabilities $p \in [0, 1]^E$, revenue $\lambda \in \mathbb{Q}_+^K$.
procedure LOCALSEARCH(I, p , REVENUE)
 $S \leftarrow \{e \in E: p(e) > 0\}$
 $\lambda^* \leftarrow \text{REVENUE}(I, p)$
while $\Delta > \Delta_0$ **do**
 Choose $E^+, E^- \subset S$ with $E^+ \cap E^- = \emptyset$ and $|E^+|, |E^-| \leq k$.
 $\Delta' \leftarrow \min\{\Delta, \sum_{e \in E^-} p_e, \sum_{e \in E^+} 1 - p_e\}$
 $p'_e \leftarrow p_e$ for all $e \in E \setminus (E^+ \cup E^-)$
 while $\exists e \in E^-$ with $p_e < \Delta' / |E^-|$ **do**
 $p'_e \leftarrow 0; E^- \leftarrow E^- \setminus \{e\}; \Delta' \leftarrow \Delta' - p_e$
 $p'_e \leftarrow p_e - \Delta' / |E^-|$ for all $e \in E^-$
 $\Delta' \leftarrow \min\{\Delta, \sum_{e \in E^-} p'_e, \sum_{e \in E^+} 1 - p'_e\}$
 while $\exists e \in E^+$ with $p_e > 1 - \Delta' / |E^+|$ **do**
 $p'_e \leftarrow 1; E^+ \leftarrow E^+ \setminus \{e\}; \Delta' \leftarrow \Delta' - p_e$
 $p'_e \leftarrow p_e - \Delta' / |E^+|$ for all $e \in E^+$
 $\lambda \leftarrow \text{REVENUE}(I, p')$
 if $\lambda > \lambda^*$ **then**
 $p' \leftarrow p; \lambda^* \leftarrow \lambda$
 $\Delta \leftarrow \Delta \cdot \varepsilon$
return (p, λ^*)

The Algorithm. In addition to an instance of LMP_X^L , the input for the algorithm consists of a candidate subset $S \subseteq E$ of the graph edges and an initial setting of probabilities on these edges. An improving move of the local search chooses two disjoint subsets $E^+, E^- \subset S$ with $|E^+|, |E^-| \leq k$ for an input parameter k . The probabilities on the edges in E^- are uniformly decreased up to a total decrease of $\Delta' = \min\{\Delta, \sum_{e \in E^-} p_e, \sum_{e \in E^+} 1 - p_e\}$, where $\Delta > 0$ is an exponentially decreasing step size. Then, the probabilities on the edges in E^+ are uniformly increased up to a total increase of Δ' . The framework then recomputes the followers' response, and accepts the move, if the leader's profit increases or reverts it, otherwise. The algorithm terminates if the improvement in the objective is below a given threshold for a given number of consecutive iterations.

Initial Solutions. The performance of the local search framework depends crucially on the choice of candidate edges and the corresponding initial solution. We tested several methods for generating such solutions. The first is solving the LP relaxation [LP] and using the support of the resulting solutions as candidate edge set. The second is computing a minimum cardinality directed multicut in the graph separating all terminal pairs, and then distributing the budget uniformly among the arcs in the cut. In addition, we also used the solutions from the MIP formulation given by Borndörfer et al. (2013) for the fixed price setting.

4.5. Elastic Demands

The models discussed so far assume that passenger demand is a fixed quantity, independent of the decisions of the leader. In reality, however, passengers are usually not confined to use public transit and might resort to an outside option if costs for traveling within the transit network are too high. In this section we show how to incorporate such *elastic demands* (Cole et al. 2006) into our model.

To extend our model, we suppose there is a value $H_i > 0$ for each commodity $i \in K$ indicating the total cost of using a different mode of transportation, e.g., a car. A passenger of commodity i will make use of this outside option whenever the cheapest option to travel within the transit network exceeds that of the outside option, i.e., if $\min_{P \in \mathcal{P}_i} f_{X,p,i}(P) > H_i$. In particular, we obtain the following revenue functions for the fixed and flexible fare variants of our model with elastic demands, respectively:

$$\Gamma_{X,i}^{\text{fix,elastic}}(p) := \begin{cases} \Gamma_{X,i}^{\text{fix}}(p) & \text{if } \min_{P \in \mathcal{P}_i} f_{X,p,i}(P) \leq H_i \\ 0 & \text{otherwise} \end{cases} \quad \text{and}$$

$$\Gamma_{X,i}^{\text{flex,elastic}}(p) := \min\{\Gamma_{X,i}^{\text{flex}}(p), H_i - \text{SP}_c(s_i, t_i)\}.$$

Note that we assume the cost of the outside option to be constant, i.e., the underlying utility function of a passenger exhibits the *constant reservation utility structure* (Acemoglu and Ozdaglar 2007). At the end of the paper, we comment on the possibility of nonconstant reservation utilities that take into account congestion effects occurring for the outside option. We now show that all results presented so far still hold in the presence of elastic demands.

First note that adding elastic demands does not affect the complexity of the followers' problem: It is still sufficient to compute the cheapest fare evasion option, i.e., a path $P \in \mathcal{P}_i$ minimizing $f_{X,p,i}(P)$, which is then compared with the option of buying a ticket and, now in addition, to the cost of the outside option H_i .

For the leader's problem, we can add a special path P_i^H to the set of paths \mathcal{P}_i in the same fashion as we used P_i^T to describe the option of buying a ticket in Section 2.2. In this case, we define $f_{X,p,i}(P_i^H) := H_i$ and $\pi(P_i^H) := 1$. With this change, the proofs for Lemma 12 and Theorem 8 remain intact. Therefore, OPT_{LP} still constitutes an upper bound for all variants of our model and the approximation guarantees for the flexible fare variants are preserved. The same holds for the exact formulations discussed in Section 4.3. Finally, since the complexity of computing the followers' response does not change, the local search procedure can easily be adapted to the elastic demand setting by introducing the additional special path.

5. Computational Study

In this section we present an extensive computational study on a broad set of realistic instances, assessing the solution quality of our local search approach and the impact of our new modeling approach compared to existing models. The source code of our algorithm and the test instances are available as part of the online companion to this paper.

5.1. Test Instances

Real-World Networks. Experiments were performed on two different instance sets. The first set comprises the complete networks of the Amsterdam subway system and the Dutch railway system, as well as a sub-network of the latter restricted to major transport hubs. The data was acquired from Dutch Railways (2014). Due to privacy regulations, no real-world passenger data was available. Therefore, for each network, we generated 10 instances each with 25, 50, 100, and 200 commodities by choosing pairs of vertices uniformly at random and drawing the corresponding demands uniformly at random from the interval $[1, 50]$. The costs representing travel time (in monetary units) were calculated from actual transit times with a conversion rate of 0.132 euro per minute travel time (Dutch Railways 2014).

Randomly Generated Instances. The second instance set comprises randomly generated planar graphs exhibiting characteristics similar to those of real-world networks (Von Ferber et al. 2009). The graphs were generated using an approach similar to the one described by Denise et al. (1996). Vertices are distributed uniformly at random in the plane; iteratively, a vertex is chosen uniformly at random and connected to the closest neighbor that can be connected without violating planarity. This is repeated $3|V| - 6$ times, after which disjoint connected components are connected using nearest Euclidean neighbours to ensure that the entire graph is connected. All arcs are present in both directions, as common in transit networks. For the randomly generated instances, ticket prices were set to

$$T_i = b + m \cdot \frac{SP_c(s_i, t_i)}{\max_{v, w \in V} SP_c(v, w)},$$

where b is a base price and $b + m$ the maximum ticket price allowed in the network. This linear formula is a simplification of the formula used for official regulations regarding ticket prices in public transport networks (Gijsbers 2013, Dutch Railways 2014). We generated 10 graphs for each possible combination of $|V|, |K| \in \{25, 50, 100, 200\}$.

For all instances, we tested 20 different values of budgets from the range of 0.2 to 25. This setup yields 800 instances for each of three real-world networks and each of the four graph size classes of the randomized

set, leading to a total of 5,600 instances. Our study consists of seven graph sets, three based on real-world transit networks and four generated using the randomization procedure described above. Each set contains 40 different combinations of graphs, customers, and demands. For each of these combinations, 20 different budget values were tested, leading to 800 different instances in each graph set; see Table 3 in Section 5.4 for average graph sizes.

5.2. Computational Setup

Algorithms. As algorithms for computing start solutions of the local search, we tested the LP relaxation (LP) and a minimum cardinality multicut computed using a standard MIP formulation (MC). To assess the impact of the more precise followers' objective function in our model as compared to existing approaches, we additionally computed the mixed integer programming solutions from Borndörfer et al. (2013) (MIP). The followers' response in the local search procedure was computed using the exact algorithms for the respective variants presented in Section 3. After initial experiments for fine-tuning the parameters of the framework, it turned out that restricting to $k = 1$, i.e., probability shifts from one edge to another, is already sufficient for obtaining close-to-optimal solutions within 30 iterations. We also set the initial step length Δ to 0.1, decreasing it by a factor of 0.9 in every iteration.

Implementation Details. All algorithms have been implemented in Java and compiled using jre7 on Windows 7 Enterprise. Computations have been performed on a machine with Intel Core 2 Duo CPU (GHz, 64 bit) and 4 GB of memory using CPLEX 12.4 API for Java for the mathematical programs.

5.3. Results

Solution Quality. Table 1 and Figure 1 show average gaps for all models on the test instances. With the exception of few instances with very high budget, results from LP^{LS} consistently dominated those of MC^{LS}. Therefore, we omitted stating the results of the latter. Solutions for the fixed-fare variant are within 95% of the upper bounds on average, while solutions for the flexible-fare variant are within 97.5%. This slight difference can be explained by the fact that the same upper bound was used for all four variants.

Computation Times. Table 2 shows average computation times for the various algorithms. Using the LP for computing start solutions, a local optimum was reached within less than a minute for most instances, with the only exception being very large graphs combined with the (nonpolynomial) nonadaptive followers' response. Furthermore, note that for large instances with many commodities, and for budgets higher than 6, the mixed integer programming formulation could not be solved in reasonable time.

Table 1. Average ratios between solutions and upper bounds in percent.

Graph set	LMP _N ^{fix}				LMP _A ^{fix}			LMP _N ^{flex}		LMP _A ^{flex}	
	Best	LP	LP ^{LS}	Δ_{MIP}	Best	LP	LP ^{LS}	LP	LP ^{LS}	LP	LP ^{LS}
Nlmajor	97.4	94.8	97.3	2.09	97.4	94.1	97.2	96.4	98.0	96.4	98.0
Nlcomplete	94.6	91.6	93.2	1.28	93.8	90.9	92.6	97.8	97.9	97.7	97.9
Adammetro	98.2	94.7	97.7	7.09	98.2	94.7	97.8	94.7	97.8	94.7	97.8
Small	97.0	92.1	96.3	4.54	96.8	88.1	96.0	95.3	97.5	95.3	97.4
Medium	96.3	90.7	95.2	4.79	96.0	88.9	95.0	95.5	97.5	95.4	97.3
Large	95.7	90.4	95.1	4.82	95.3	88.6	94.5	96.6	98.0	96.5	97.8
Huge	95.6	89.1	94.1	4.03	94.6	85.2	92.7	96.5	97.9	96.4	97.7

Notes. “Best” denotes the average over the best solutions found for each instance, LP denotes the solutions found by [LP], LP^{LS} denotes the solution found by performing the local search heuristic on the LP solution, and Δ_{MIP} denotes the improvement of “best” compared to the solution found by using the MIP of Borndörfer et al. (2013).

Table 2. Average computation time in seconds.

Graph set	LP	LP _N ^{LS}	LP _A ^{LS}	MIP	MIP _N ^{LS}	MIP _A ^{LS}
Nlmajor	0.11	20.3	3.12	47.9	20.4	3.97
Nlcomplete	1.69	100	9.3	651	98	9.3
Adammetro	0.17	13.7	3.0	1.59	11.8	2.5
Small	0.15	3.2	0.34	18.7	3.55	0.38
Medium	0.18	9.2	0.85	30.6	10.0	0.93
Large	0.47	37.9	3.16	190	39.0	3.35
Huge	1.7	91.8	7.98	700	102.3	8.83


Notes. LP and MIP denote the solutions found by the corresponding mathematical program, and Alg_X^{LS} with Alg ∈ {LP, MIP} and X ∈ {A, N} denotes the solution found by performing the local search heuristic on the solution found by algorithm Alg, using the followers’ response X.

Impact of the Budget Size. We also investigated the impact of budget sizes on the optimality gap and the achieved profit; a visualization is shown in Figure 2. In the range of $B \in [1, 4]$, the ratio of obtained profit and upper bound exhibits a bathtub curve behavior until

it stabilizes at a ratio of 1 for larger budgets (which is expected as for a large enough budget, all passengers can be forced to buy a ticket). The achievable profit as a function of the budget is concave in all investigated examples, which might be a universal property of the optimal value function (as a function of the budget).

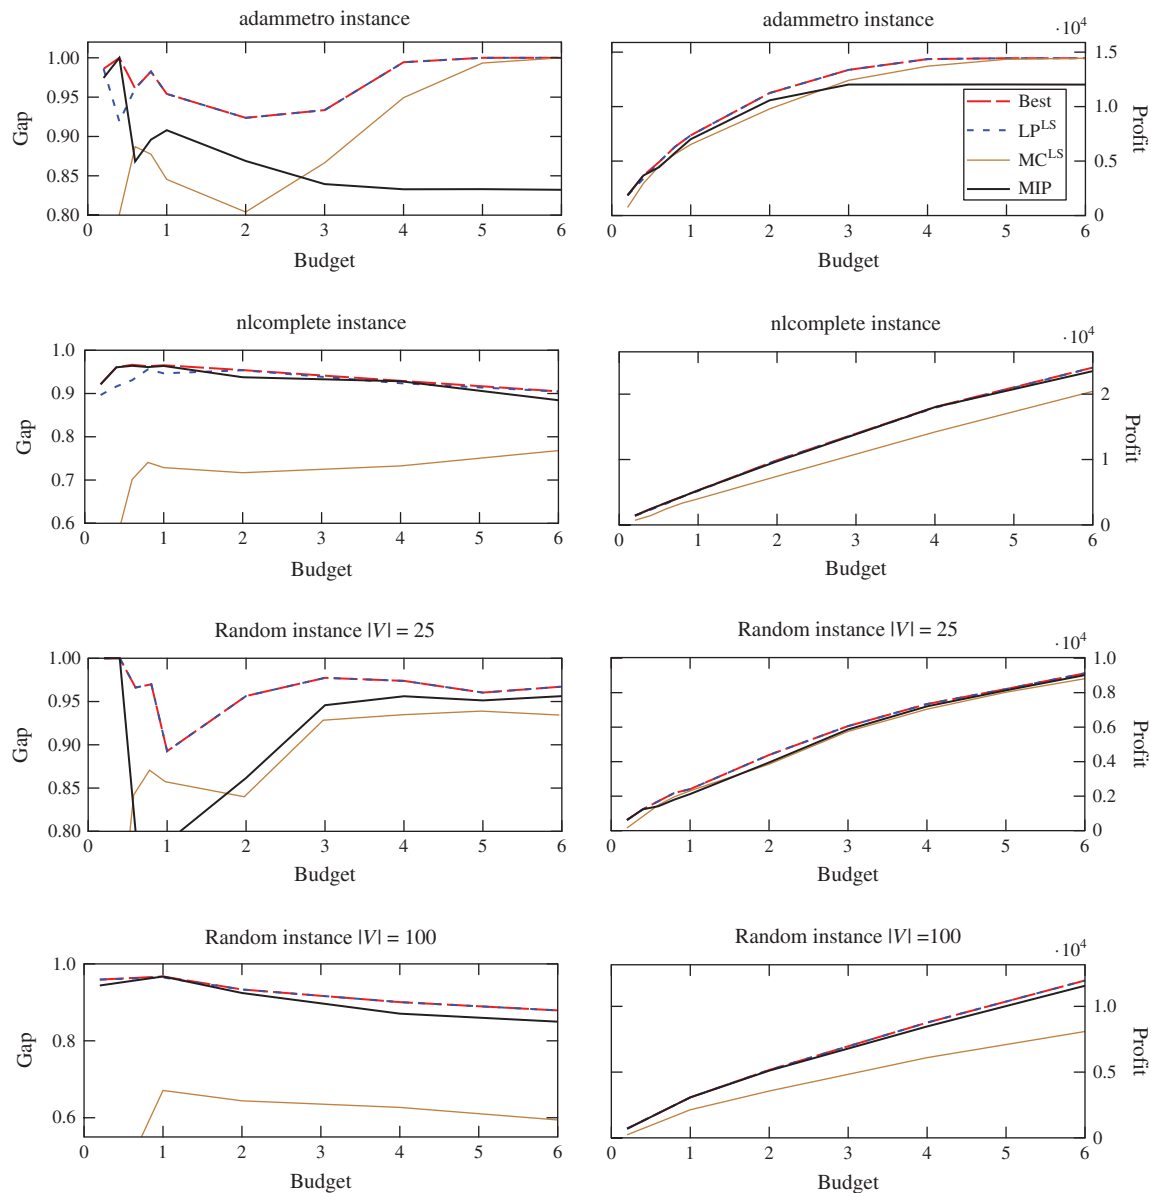
Comparison with Existing Models. For the fixed-fare variant with nonadaptive followers, we can compare our modeling approach to that of Borndörfer et al. (2013). We solved the MIP formulation proposed in Borndörfer et al. (2013) and computed the leader’s profit resulting from the realistic response of followers (i.e., without linearization of their objective function). Comparing these solutions to the ones derived from our local search procedure, we observed that our approach yields an increase in profit of about 5% on average on the randomly generated instances, about 7.5% on the metro, and about 2% on the railway network. In fact, for some of the metro instances, the increase exceeds 20%; see Table 1 and Figure 1.

Figure 1. (Color online) Ratio of profits to upper bounds.



Notes. The diagrams show the distribution of the ratios for instances in the respective test sets. The mark inside each box denotes the median, boxes represent lower and upper quartiles, and the whisker ends show the minimum and maximum, respectively, apart from possible outliers marked by a cycle.

Downloaded from informs.org by [200.89.69.180] on 03 February 2017, at 06:51. For personal use only, all rights reserved.

Figure 2. (Color online) Graphs illustrating the influence of the budget on the solution quality.

Notes. The graphs depict gaps and profits in dependence of the budget for the adammetro instance, the nlcomplete instance, a random instance with 25 vertices and a random instance with 100 vertices, each with 50 commodities for the LMP_N^{fix} model. Note that the legend in the top-right graph applies to all graphs. Graphs on the left show gaps relative to the upper bound, graphs on the right show the corresponding profit. “Best” denotes the best solution found, LP^{LS} and MC^{LS} denote the solutions found by performing the local search heuristic on the LP solution resp. MultiCut solution, and MIP denotes the solution found by the corresponding mathematical program of Borndörfer et al. (2013).

5.4. Computation of Exact Solutions

In the following we present some (limited) results regarding exact approaches for the four variants of the problem.

Globally Optimal Solutions Using BARON. To find globally optimal solutions for small instances, we use the exact branch-and-bound-based nonlinear mixed integer global optimization solver BARON (Tawarmalani and Sahinidis 2005) to solve the corresponding nonlinear formulations $[NLP_A^{\text{flex}}]$, $[NLP_A^{\text{fix}}]$, $[MINLP_N^{\text{flex}}]$, and $[MINLP_N^{\text{fix}}]$ described in Section 4.3.

Note that all these formulations are nonconvex and therefore not easily solvable by (standard) mixed integer *convex* programming solvers. Secondly, for $[MINLP_N^{\text{flex}}]$, the modeling step of moving to a path formulation comes with a possibly exponential growth in the number of decision variables, since the number of paths can be exponential in the size of the network. To reduce this possibly large number, for every commodity $i \in K$, we can prune all paths $P \in \mathcal{P}_i$ for which $c(P) \geq SP_c(s_i, t_i) + T_i$, since these paths will never be chosen in an optimal solution. Furthermore, calculating

Table 3. Average instance sizes and gaps to exact solutions.

Graph set	$ V $	$ E $	T_i	F	$P_{[\text{MINLP}_N^L]}$	$\text{GAP}_N^{\text{flex}}$	$\text{GAP}_A^{\text{flex}}$	$\text{GAP}_N^{\text{fix}}$	$\text{GAP}_A^{\text{fix}}$
Nlmajor	23	60	219	423	342	98.7	98.7	98.1	98.0
Nlcomplete	341	864	753	1,630	—	(97.8)	(97.7)	(94.6)	(93.8)
Adammetro	45	88	289	591	25	99.2	99.6	98.8	98.3
Small	25	95	218	440	23,438	97.6	98.0	(97.0)	(96.8)
Medium	50	195	302	623	—	(96.3)	(96.0)	(97.5)	(97.3)
Large	100	399	404	881	—	(95.7)	(95.3)	(98.0)	(97.8)
Huge	200	806	581	1,248	—	(95.6)	(94.6)	(97.9)	(97.7)

Notes. We report average values for instances with 25 commodities only. $P_{[\text{MINLP}_N^L]}$ provides the number of path variables created for $[\text{MINLP}_N^L]$ after pruning. A dash indicates that the instances in the set were too large in terms of number of paths to fit into memory. GAP_X^L denotes the gap between our solutions for LMP_X^L and the minimum of our upper bounds and the results from the BARON computations. All values in brackets in the right set of columns indicate known gaps for those instances for which BARON could not get improved values due to memory and/or time limitations.

a minimum multicut provides a lower bound on the budget for which we obtain maximum possible revenue: setting $p_e = 1$ for every edge e in the multicut guarantees the maximum obtainable profit of $\sum_{i \in K} T_i$.

Computations were performed using BARON v.15.6.5 on the same machine as the other algorithms, and interaction with the BARON software was performed using Java. The solver was allowed 600 seconds to process each instance. If an optimal solution could not be found, the upper bound found by the solver up to that point was compared to our upper bound, and used to compare our results if this gave an improvement. The results are listed in Table 3, where the gaps for instances which were not solved using BARON are provided between brackets for completeness.

In Table 3 the average instance sizes for each graph set are presented in the left set of columns, as well as the number of path variables in $[\text{MINLP}_N^L]$ after pruning. Even though the network sizes of nlmajor, ADAMmetro and the small random graphs are relatively the similar in the number of arcs and nodes, the number of paths after pruning shows a huge difference. This difference is easily explained by observing the structure of the graphs. Both small real-life networks are star-like in structure: There is a single center of clustered nodes (forming an almost-clique), and terminals are connected to the center via possibly parallel paths. On the other hand, for the randomly generated graphs there is no single center, but nodes are evenly distributed, e.g., resembling the graph structure of the network of the Deutsche Bahn. It turned out that, using BARON, we were not able to compute any exact solution or useful upper bounds for the random instances using our exact formulations and the time limits as stated above.

For the graph sets nlmajor and ADAMmetro, we obtained exact solutions for many instances, and in some cases improved upper bounds as shown in Table 3. Comparing our results with the solutions provided by BARON shows that the solutions found

by the local search algorithm for instances derived from nlmajor and ADAMmetro are consistently near-optimal (within 2% of optimality). This can be seen as an indication that, in general, our solutions are even slightly closer to the optimum than suggested by the data in Table 1.

6. Conclusions

In this paper we introduced and studied models for the optimization of fare inspection strategies in transit systems taking into account realistic passenger behavior. We developed efficient algorithms for the passengers' reaction (given a distribution of control probabilities) as well as for the overall bilevel optimization problem of computing control probabilities and ticket prices. We demonstrated in an extensive computational involving a total of 5,600 instances generated from real-world networks as well as randomly generated instances that our algorithmic approach leads to high quality solutions that are within 5% of the calculated upper bounds on average.

From a theoretical and practical point of view, our work raises several interesting questions:

Complexity of the nonadaptive followers' minimization problem. The nonadaptive followers' problem FMP_N discussed in Section 3 constitutes a natural generalization of the classic shortest path problem, which is of interest beyond the concrete application in the present work. While the quasi-polynomial bound on the number of breakpoints for the parametric shortest path problem used in this paper is tight, a polynomial time algorithm for FMP_N still could be achieved using a different approach.

Complexity of the parametric shortest path problem in planar graphs. As most real-world transit networks are based on planar or almost planar graphs, finding efficient algorithms for this type of network is an important task of research in this area. In this context, Nikolova's (2009) conjecture on the polynomiality of

the parametric shortest path problem in planar graphs is of particular interest, as it implies polynomial time algorithms for a broad class of shortest path problems in planar graphs, including FMP_N . While we proved the conjecture for the special case of series-parallel graphs, the case of arbitrary planar graphs remains open and is an interesting subject of future research.

Approximation complexity of the leader's problem.

We have shown strong NP-hardness of all variants of the leader's optimization problems, even in restricted special cases. It remains a challenging open problem to either improve upon our approximation guarantees, or to prove APX-hardness results closing the gap with our approximation guarantees.

Robustness against suboptimal followers' reaction.

If passengers only have incomplete information about the distribution of control probabilities, they might respond suboptimally. Thus, it would be interesting to derive more general models that yield solutions that are robust to suboptimal behavior.

Operational planning of inspection routes. In the model presented in this paper, average control frequencies are specified by assigning control probabilities to the edges. This is a suitable assumption in the context of tactical planning of network control, i.e., over a longer time horizon. It would be interesting to also consider a more fine-grained model for operational day-to-day planning that involves the dynamic route planning of inspectors moving at different points in time to different checking spots.

Multimodal transportation model. In the model described in Section 4.5, restricting reservation utilities to be constant assumes that alternative modes of transportation are not affected, e.g., by congestion. Removing this restriction gives rise to a bilevel multimodal transportation model, about which not much is known. Taking into account that this setting is not uncommon in real-life applications, it would be interesting to further investigate these types of models.

Acknowledgments

The authors thank Ralf Borndörfer and Elmar Swarat for giving an inspiring talk on the topic of toll enforcement that motivated this study. The authors further thank Evdokia Nikolova for pointing out the connection to parametric shortest paths, which helped to considerably simplify and generalize the results in Section 3.

References

Acemoglu D, Ozdaglar AE (2007) Competition and efficiency in congested markets. *Math. Oper. Res.* 32(1):1–31.
 Borndörfer R, Sagnol G, Swarat E (2012b) A case study on optimizing toll enforcements on motorways. Ravizza S, Holborn PL, eds. *Proc. 3rd Student Conf. Oper. Res.* Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany), 1–10.

Borndörfer R, Omont B, Sagnol G, Swarat E (2012a) A Stackelberg game to optimize the distribution of controls in transportation networks. *Proc. 3rd Internat. Conf. Game Theory for Networks* (Springer, Berlin), 236–247.
 Borndörfer R, Buwaya J, Sagnol G, Swarat E (2013) Optimizing toll enforcement in transportation networks: A game-theoretic approach. *Electronic Notes in Discrete Math.* 41:253–260.
 Bouhtou M, Grigoriev A, van Hoesel S, van der Kraaij AF, Spiessma FCR, Uetz M (2007) Pricing bridges to cross a river. *Naval Res. Logist.* 54(4):411–420.
 Briest P, Hoefer M, Krysta P (2012) Stackelberg network pricing games. *Algorithmica* 62(3–4):733–753.
 Brotcorne L, Labbé M, Marcotte P, Savard G (2001) A bilevel model for toll optimization on a multicommodity transportation network. *Transportation Sci.* 35(4):345–358.
 Brotcorne L, Labbé M, Marcotte P, Savard G (2008) Joint design and pricing on a network. *Oper. Res.* 56(5):1104–1115.
 Carstensen PJ (1983) The complexity of some problems in parametric linear and combinatorial programming. Ph.D. thesis, University of Michigan.
 Chilean Subsecretary of Transportation (2014) Fare evasion in Transantiago. <http://www.fiscalizacion.cl/index.php/evasion-de-transantiago-alcanza-192-y-registra-nivel-mas-bajo-desde-diciembre-de-2010/>.
 Cole R, Dodis Y, Roughgarden T (2006) Bottleneck links, variable demand, and the tragedy of the commons. *Proc. Seventeenth Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '06* (ACM, New York), 668–677.
 Denise A, Vasconcellos M, Welsh DJA (1996) The random planar graph. *Congressus Numerantium* 113:61–79.
 Dijkstra EW (1959) A note on two problems in connexion with graphs. *Numerische Mathematik* 1(1):269–271.
 Dutch Railways (2014) NS tarieven (Dutch tariff regulations). <http://www.ns.nl/reizigers/klantenservice/klantenservice/tarieven>.
 Gijssbers BCM (2013) Directeur Openbaar Vervoer en Spoor, Private communication.
 Gómez-Lobo A (2012) The ups and downs of a public transport reform: The case of Transantiago. Technical report, Universidad de Chile, Santiago, Chile.
 Gondran M, Minoux M, Vajda S (1984) *Graphs and Algorithms* (John Wiley & Sons, New York).
 Henzinger MR, Klein P, Rao S, Subramanian S (1997) Faster shortest-path algorithms for planar graphs. *J. Comput. System Sci.* 55(1):3–23.
 Jain M, Conitzer V, Tambe M (2013) Security scheduling for real-world networks. *Proc. 2013 Internat. Conf. Autonomous Agents and Multi-agent Systems* (IFAAMAS, Richland, SC), 215–222.
 Klinz B, Woeginger GJ (2004) Minimum-cost dynamic flows: The series-parallel case. *Networks* 43(3):153–162.
 Labbé M, Marcotte P, Savard G (1998) A bilevel model of taxation and its application to optimal highway pricing. *Management Sci.* 44(12):1608–1622.
 Lin KY, Atkinson MP, Chung TH, Glazebrook KD (2013) A graph patrol problem with random attack times. *Oper. Res.* 61(3):694–710.
 Lorenz DH, Raz D (2001) A simple efficient approximation scheme for the restricted shortest path problem. *Oper. Res. Lett.* 28(5):213–219.
 Nikolova E, Kelner JA, Brand M, Mitzenmacher M (2006) Stochastic shortest paths via quasi-convex maximization. Azar Y, Erlebach T, eds. *Proc. 14th Annual European Sympos. Algorithms, ESA '06* (Springer, Berlin), 552–563.
 Nikolova EV (2009) Strategic algorithms. Ph.D. thesis, Massachusetts Institute of Technology.
 Roch S, Savard G, Marcotte P (2005) An approximation algorithm for Stackelberg network pricing. *Networks* 46(1):57–67.
 Tambe M (2012) *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (Cambridge University Press, New York).

- Tawarmalani M, Sahinidis NV (2005) A polyhedral branch-and-cut approach to global optimization. *Math. Programming* 103(2): 225–249.
- Torres-Montoya M (2014) Tackling fare evasion in Transantiago: An integrated approach. *Transportation Res. Board 93rd Annual Meeting* (Transportation Research Board, Washington, DC).
- Transport for London (2010) Bus fare evasion communications development. Technical report, Transport for London. <http://www.tfl.gov.uk/assets/downloads/customer-research/bus-fare-evasion-communications-development-presentation.pdf>.
- Von Ferber C, Holovatch T, Holovatch Y, Palchykov V (2009) Public transport networks: Empirical analysis and modeling. *Eur. Physical J. B* 68(2):261–275.
- Washburn A, Wood K (1995) Two-person zero-sum games for network interdiction. *Oper. Res.* 43(2):243–251.
- Wayne KD (1999) Generalized maximum flow algorithms. Ph.D. thesis, Cornell University, Ithaca, NY.
- Yin Z, Jiang AX, Johnson MP, Kiekintveld C, Leyton-Brown K, Sandholm T, Tambe M, Sullivan JP (2012) TRUSTS: Scheduling randomized patrols for fare inspection in transit systems. Fromherz MPJ, Munoz-Avila H, eds. *Proc. 24th Conf. Innovative Applications of Artificial Intelligence* (AAAI, Palo Alto, CA), 2348–2355.

José Correa is professor of operations research in the Department of Industrial Engineering at the Universidad de Chile. His current research deals with game theory, pricing, and mechanism design in operational contexts including decentralized networks, scheduling, and inventory clearing.

Tobias Harks is professor of optimization in the Institute of Mathematics at the University of Augsburg. His research interests cover the design of algorithms, algorithmic game theory, and discrete and continuous optimization.

Vincent J. C. Kreuzen is a lecturer in the Department of Quantitative Economics at Maastricht University. The present paper is part of his doctoral thesis on the design of approximation algorithms in operations research.

Jannik Matuschke is junior research group leader at the Chair of Operations Research at the TUM School of Management. His research interests cover combinatorial optimization and graph algorithms, in particular network optimization and its applications in operations research.